

Ou: Automating the Parallelization of Zero-Knowledge Protocols

Yuyang Sang*
Yale University

yuyang.sang@yale.edu

Ning Luo*†
Northwestern University

ning.luo@northwestern.edu

Samuel Judson
Yale University

samuel.judson@yale.edu

Ben Chaimberg
Yale University

ben.chaimberg@yale.edu

Timos Antonopoulos
Yale University

timos.antonopoulos@yale.edu

Xiao Wang
Northwestern University

wangxiao1254@gmail.com

Ruzica Piskac
Yale University

ruzica.piskac@yale.edu

Zhong Shao
Yale University

zhong.shao@yale.edu

Abstract

A zero-knowledge proof (ZKP) is a powerful cryptographic primitive used in many decentralized or privacy-focused applications. However, the high overhead of ZKPs can restrict their practical applicability. We design a programming language, *Ou*, aimed at easing the programmer’s burden when writing efficient ZKPs, and a compiler framework, *Lian*, that automates the analysis and distribution of statements to a computing cluster. *Lian* uses programming language semantics, formal methods, and combinatorial optimization to automatically partition an *Ou* program into efficiently sized chunks for parallel ZK-proving and/or verification.

We contribute:

1. A front-end language where users can write proof statements as imperative programs in a familiar syntax;
2. A compiler architecture and implementation that automatically analyzes the program and compiles it into an optimized IR that can be lifted to a variety of ZKP constructions; and
3. A cutting algorithm, based on Pseudo-Boolean optimization and Integer Linear Programming, that reorders instructions and then partitions the program into efficiently sized chunks for parallel evaluation and efficient state reconciliation.

1 Introduction

Zero-knowledge proofs (ZKPs) [GMW86] enable a prover to convince a verifier about the truth of some statement without revealing why. Recent advances by the cryptographic research community have brought a tremendous improvement in the efficiency of general-purpose zero-knowledge (ZK) proving (see [zpk] for a comprehensive survey), as well as numerous creative applications of ZKPs (e.g., [BCG⁺14, BMRS20, BCG⁺20, FPS⁺18, ACBM08]). However, various practical challenges make the adoption of ZK unfeasible for tasks of realistic sizes. A particular problem is the lack of a comprehensive toolkit of programming languages and compilers providing a full suite of effective, intuitive, *parallel-aware*, and general methods for writing complex statements intended for efficient ZKPs.

*These authors contributed equally to this work.

†Part of this work was completed while the author was at Yale University.

Challenge 1: Simple and scalable development. One significant challenge for applying ZKPs to new problems is describing the statement to be proven in a way that is at the same time natural to the programmer and conducive to compilation.

Most existing languages and compilers for ZK (and for the closely related problem of secure multi-party computation, or MPC) support only the ‘circuit model’ based on *data-oblivious* computation, where control flow must be independent of all private variables [BLW08, KsMB13, LWN⁺15]. They encode the statement as a monolithic program and emit an appropriate circuit representation for their target protocol. Although this model is theoretically complete [PF79, GO96], it does not capture the full flexibility and power of modern ZKP constructions. For example, many statements can be accelerated by non-determinism (*e.g.*, [BCG⁺13, BCTV14]), where the prover provides extra hints to the computation. Other statements (*e.g.*, 3-coloring and Hamiltonian cycles in graphs [Gol09]) can be verified with high concrete efficiency under a probabilistic guarantee: the verifier provides extra random challenges, to which the prover replies with challenge-dependent hints.

We contribute a more general-purpose ZK programming framework able to support proofs over multiple computational models. Our aim is to ease development for new domains – especially so for ZK statements that are more naturally expressed in those alternative models, or for which those models provide superior performance.

Challenge 2: Scalability in hardware resources. All sufficiently advanced technology may be indistinguishable from magic [Cla13], but the ‘magic’ of ZK often depends as much in practice on massive hardware resources as on mathematical ingenuity. These demands lead to practical bottlenecks in computation and/or communication resources when scaling ZKPs to large statements. For example, most zkSNARK protocols [Gro10, BCC⁺16, BBB⁺18, WTs⁺18, Set20] require temporary storage linear in the running time of the computation being proved. Although VOLE-based ZKPs [WYKW21, BMRS21, DIO20] use less memory, they only do so at the expense of a significant increase in required bandwidth.

To aid practicality, a line of recent work has started to “scale out” ZK protocols: instead of assuming the prover and/or verifier have access to one giant machine, they assume access to a cluster of machines and the capability to distribute the task amongst these resources. For example DIZK [WZC⁺18] distributes the prover computation of Groth’s ZK protocol [Gro16] by manually partitioning the R1CS constraints into equal-sized chunks. zkBridge [XZC⁺22] can distribute the prover computation of Virgo [ZXZS20] efficiently to multiple servers when the statement already has a high degree of parallelism. Meanwhile, EZEE [YHKD22] distributes among multiple machines the prover-verifier computation and communication of an interactive, garbled-circuit-based zero-knowledge proof protocol [JKO13]. Giraffe [WJB⁺17] addresses the challenge of automatically distributing computations in the verifiable outsourcing context by breaking down overly large computations that cannot be outsourced as a whole.

Though efficient in deployment, in order to distribute their chosen ZK protocol over a cluster of servers each of these works require the developer to first both i) modify the ZKP protocols to be amenable to distributed computation; and ii) to manually partition the computation into smaller chunks to be proven in parallel. Such manual processes may be error-prone and lead to suboptimal solutions, especially for statements of real-world complexity.

1.1 Summary of Contributions

In this paper, we design a programming framework – a language `Ou` and a supporting compiler architecture `Lian` – enabling programmers to write zero-knowledge statements without thinking about distributed computation. The compiler automatically and efficiently chunks an `Ou` program for efficient parallel proving or verification.

1. Our framework provides a C-like programming language with annotations so that developers, including those without deep knowledge of cryptography, can easily write ZK applications even when using

protocols with advanced features.

2. We develop a compiler and static analysis pipeline, written in OCaml, that calculates the local and communication costs of each program instruction, then automatically finds an efficient way to chunk the corresponding ZK statement. The optimality measure minimizes the sum of the maximum per-machine cost and cumulative communication costs. This analysis can distribute ZK statements incorporating randomized verification.
3. We prove our system sound, *i.e.*, we show the semantics of the distributed verification of substatements is equivalent to verifying the whole statement centrally.
4. Empirically, we show that for many statements, including gradient descent and Merkle Tree, our framework can automatically find efficient partitioning of the statements, thus achieving speedup proportional to the number of machines available.

What this paper is not about. Lian does not establish a program’s functional correctness or data-oblivious nature, nor the hardness of witness finding for the ZK statement. We assume a program written in our language is sound and secure for its intended ZK application, and is suitably constructed for the target protocol: *e.g.*, we assume the program is data-oblivious if the target is circuit-based.

1.2 Technical Challenges and Solutions

Ou and Lian address a series of challenges arising from each of language design, cryptography, and combinatorial optimization.

Cryptographic insight. To enable arbitrary partitioning of a statement, the underlying ZK protocol needs to be ‘flexible’: it must allow proving multiple statements of the same witness while ensuring consistency. To this end, our framework focuses on the commit-and-prove paradigm. The prover first commits to their private inputs, and then constructs a proof establishing some (public) relationship among the committed values without revealing them to the verifier. Because the inputs are committed the prover can establish multiple statements on them, all while the verifier is certain their values remain consistent. Any ZKP system can support this capability by proving the consistency of commitments in ZK, but certain protocols have direct support which incurs less cost (*e.g.*, [WYKW21, BMRS21, DILO22, JKO13, CFQ19, CFF⁺21, HK20, AGM18]). Given such a ZK backend secure under parallel composition, regardless of its internals, *e.g.*, the use of Fiat-Shamir or random oracle, we can then distribute the proof of any ZK statement in parallel.

During compilation, we first partition the statement into multiple substatements, before the prover precomputes the input/output of each substatement locally. The prover can then in parallel prove to the verifier the correctness of each substatement and then establish the consistency of the input/output between substatements. Although this works in principle, designing a scalable system that automatically distributes the computation with minimum runtime- and developer-overhead is still challenging. For example, in ZK data dependency rarely constrains statement partitioning, unlike in normal computation or MPC where it does so always. So in ZK even inherently sequential statements can often be parallelized.

Automatic statement partitioning. In ZK the prover knows both the public and private inputs, and can thus ‘predict’ all intermediate values. As a result, data dependencies can be resolved before proving: computation of $y = f(g(x_1), x_2)$ can be partitioned into two parallel verification tasks, $t = g(x_1)$ and $y = f(t, x_2)$. The prover, with inputs x_1 and x_2 , just appends t to form an *extended witness*. For ordinary computation this parallelization would be infeasible, because f depends on g . The prover’s ability to ‘locally precompute’ creates new opportunities to maximize parallelism.

To partition effectively, after parsing and typechecking our architecture undertakes two distinct special compilation phases: *shallow simulation* and *deep simulation*. The first automatically partitions the ZK

```

1 #define S 100
2 #define secparam 128
3 plocal1 plc1 int[S][S] mmult_plcmx(...) {...}
4 pvt2 int[S] mmult_pvtvec(...) {...}
5 pvt2 int[S] mmult_pubvec(...) {...}
6
7 void frvlds(pvt1 int [S][S] M,
8             pvt1 int [S][S] M1, pvt1 int [S][S] M2) {
9     /* array of verifier controlled randomness */
10    pub2 int [S] s = {0};
11
12    int t = 0;
13    /* repeat for secparam number of times */
14    while (t < secparam) {
15        int i = 0;
16        while (i < S) {s[i] = v_rand(0, 1); i = i+1;}
17
18        /* compute z = (M1 * (M2 * s)) */
19        pvt2 int[S] w = mmult_pubvec(M2, s);
20        pvt2 int[S] z = mmult_pvtvec(M1, w);
21
22        /* compute q = (M * s) */
23        pvt2 int[S] q = mmult_pubvec(M, s);
24
25        /* check that q and z are the same */
26        i = 0;
27        while (i < S) {
28            assert (q[i] == z[i]); i = i+1;
29        }
30        t = t+1;
31    }
32 }
33
34 unit mmult(plc1 int [S][S] M1, plc1 int [S][S] M2) {
35     /* prover locally computes M1 * M2 = M */
36     plc1 int[S][S] M = mmult_plcmx(M1, M2);
37     return frvlds(commit(M), commit(M1), commit(M2));
38 }

```

Figure 1: Running example. We consider this Ou implementation of Freivalds’ algorithm for randomized verification of 100x100 matrix multiplication.

statement and then selects variables as witnesses (the appendices to the input, like t) for each of the κ sub-statements. The parameter κ is user-chosen, and will equal the number of available compute cores in the cluster. Deep simulation then uses the prover’s private inputs to concretely compute the values of those extended witnesses, as well as of any public variables.

Shallow simulation proceeds by using live variable analysis to compute the costs of each instruction in the statement, and then models them as a directed acyclic graph for partitioning by either *pseudo-Boolean optimization* (PBO) or *integer linear programming* (ILP). This requires loop unrolling and function inlining as flattening techniques. The resultant ZK substatements are then compiled to the chosen backend protocol implementation for distribution as programs to the participating servers: in our evaluation we use the VOLE-based EMP-toolkit [WMK16]. Lian’s deep simulation then propagates input-dependent values through the statement to concretely compute both public values and the prover’s witnesses for each substatement. To distinguish which values are available during our different simulation phases, we use *knowledge levels*: K_0 for (private input-independent) values that are known at compile time, and K_1 for (private input-dependent) values that are known at distribution time. This way, reusing a proof with different inputs only requires recomputing the deep simulation, and not the shallow simulation or partitioning, where the latter is likely to be the most computationally intensive compilation task.

Supporting randomized verification. As noted, randomized verification can be much more efficient than deterministic verification. However, supporting it during shallow simulation in particular is a challenge. With deterministic verification an oblivious program trace is known at compile time, which allows Lian to analyze the cost of each substatement easily. With randomized verification the program might branch on the at-the-time unknown randomness, preventing the cost from being calculated precisely as whether an expensive statement is entered may depend on a coin flip. Also, privacy imparts an ordering requirement on randomized verification: the verifier’s randomness used in checking a ZK statement can only be revealed after the prover commits to its inputs. This means that the partitioning needs to obey sequential ordering for values that are dependent on dynamic choices of the verifier.

Ou uses a third knowledge level, K_2 , to indicate that the value of a variable can only be known at runtime. The type system enforces non-interference with respect to a security lattice, allowing K_0 - and K_1 -leveled variables to flow into K_2 -leveled variables, but not the other way around. The partitioning then treats the K_2 -leveled variables as one contiguous block in order to estimate worst-case cost and enforce sequential ordering.

Handling large statements. Naively implementing the aforementioned techniques does not scale to large ZK statements. If the cost graph contains too many nodes then finding efficient chunks will be computationally infeasible. Classical graph partitioning and many related variants are NP-hard with efficient (approximate) optimization algorithms known only in certain special cases, and even then often only theoretically [BMS⁺16]. More subtly, a naive representation of each substatement, *e.g.*, as a circuit or R1CS, is necessarily linear in its running time. For large statements, especially those containing unrolled loops or inlined function calls, it is therefore crucial to work with ‘compact representations’ of substatements which are sublinear in the running time.

Ou and Lian support sublinear representations through `atomic`, a user-provided annotation to indicate that a function should be condensed, rather than flattened, during shallow simulation. By doing so, we contract the component of the cost graph corresponding to the function into a single node of cost cumulative of its constituent instructions. This reduces the size of the graph and makes optimization more practical. Even if the `atomic` function is called multiple times, only one copy of the declaration exists. This brings the partitioning problem into the realm of practicality for both PBO and ILP solving, despite its inherent hardness.

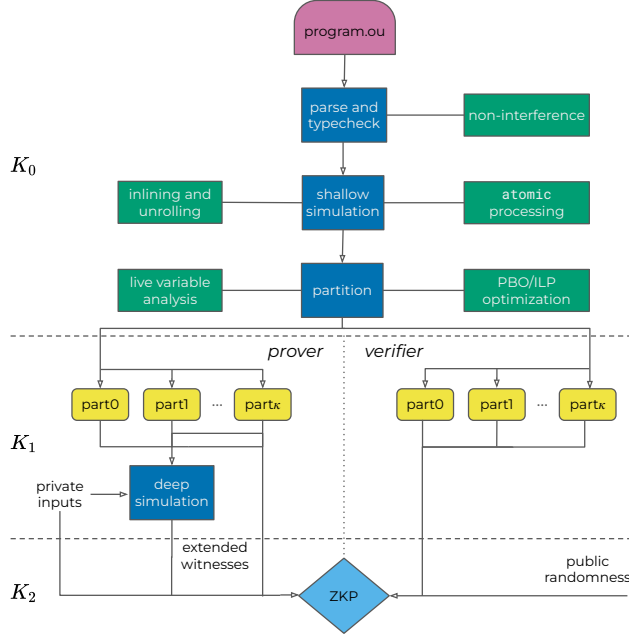


Figure 2: The Lian framework, beginning with an Ou program and ending in binaries for an implementation of a ZKP protocol. Square blocks with light text identify the compiler components: the operations (blue) and their underlying methods (green). Dark text indicates a program or protocol: the Ou program (purple semi-rounded rectangle), compiled substatement binaries (gold rounded rectangles), and the ZKP protocol itself (light blue diamond). Corresponding knowledge levels are displayed as well.

1.3 Roadmap

Sec. 2 is an overview of the Lian framework’s workflow. Sec. 3 presents the language’s syntax, typing rules, and dynamic semantics. Sec. 4 shows how the shallow simulation unfolds the program into a sequential program, and formally proves the two programs have the same behavior. Sec. 5 then describes how the compiler finds an efficient way to partition the sequential program into multiple chunks and generate a distributed program. Sec. 6 shows how the deep simulation evaluates all dependent data between chunks. It also proves that with these data, the distributed program has the same behavior as if it was run without distribution. Finally, Sec. 7 studies the effectiveness of the framework for distributing the computations of various benchmarks.

2 Lian Design Overview

Figure 2 shows the architecture and workflow of the Lian framework. The user first writes a program in Ou. Upon invocation, the compiler parses and typechecks the program. The Ou type system enforces non-interference over the security lattice given in Figure 3. This lattice specifies non-interference along two dimensions:

- i the public or private nature of the information; and
- ii the knowledge level of the information (i.e., if the information is available at compile time, distribution time or run time)

Accordingly, the lattice defines seven security domains: public $\text{pub} \times \{K_0, K_1, K_2\}$, committed and authenticated $\text{pvt} \times \{K_1, K_2\}$, and prover-local $\text{plc} \times \{K_1, K_2\}$. As syntactic sugar, we write pub0 for (pub, K_0) and similarly for the rest of the lattice. The variables in the public domain are known by both the prover and the verifier, the variables in the prover-local domain are only known by the prover, and the

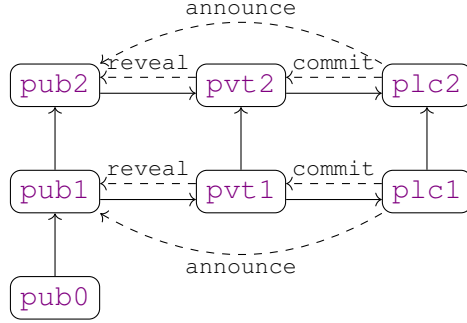


Figure 3: Security Lattice

variables in the authenticated domain are only known by the prover but the verifier holds a commitment to them. Since values known at compile time can only be public, (pvt, K_0) and (plc, K_0) do not exist.

The authenticated and prover-local variables contain information computed from the prover’s secret inputs. Public variables are usually used to guide the program’s execution, either deterministically as iterators or fixed constants, or as verifier-provided randomness. In a valid OU program every variable is annotated by the security label of a position in the lattice.¹ For example, in Figure 1 `pub0` annotates the public matrix size constant s that is known at compile time. In comparison, further up the lattice $\text{pvt2} = (\text{pvt}, K_2)$ annotates w , which depends on both the prover’s committed private values ($M2$, necessitating pvt) and on the verifier’s public randomness (s , necessitating K_2).

The security lattice enforces that the private and prover-local domains’ information never flow into the public domain except by explicit lowering through a `reveal` operation. This non-interference property [GM82, SM03] allows the safe partial execution of the program during compile time (the shallow simulation), as the control flow can be pre-determined. The inclusion of the knowledge levels in the security lattice is essential to guaranteeing that `reveal` operations and public verifier-randomness are safely handled under the non-interference guarantee.

Once typechecking completes shallow simulation commences. Shallow simulation yields a sequential program consisting of a flattened sequence of operations through loop unrolling and function inlining. With two exceptions, every instruction in the original OU program forms its own independent block in this sequential program: the constituent instructions of functions decorated by the `atomic` keyword are combined into a single block, as are all K_2 -annotated variables. In both cases, this mapping enforces they be treated as one cumulative ‘instruction’ for the purposes of partitioning. In Figure 1 for example, the suffix of the program beginning with the first `mmult_pubvec` call on L19 will be condensed, thereby guaranteeing the security of the parallel composition in the presence of public randomness.

If the user sets $\kappa = 1$, so that only a single partition is desired, then this sequential program will be directly compiled into the ZK statement as specified by the backend. However, usually $\kappa > 1$ as the developer desires parallelization, and we generally assume this to hold through the remainder of our discussion. The Lian compiler next conducts live variable analysis to track dependencies among the blocks and builds a corresponding directed acyclic graph $G = (V, E)$. An edge $(i, j) \in E$ captures that j -th instruction block depends on some information computed in the i -th instruction block. For example, in Figure 1 the `commit` (M) on L37 depends on the `mmult_plcmx` block on L36.

In order to obtain efficient partitions, the results of the live variable analysis are used to decorate G with both edge and node labels, each indicating a form of cost to the optimality of any partition. Nodes in G are labeled with their *computation cost*, which models the time required by the prover and verifier to execute the block’s proving or verification operations in the chosen ZKP protocol backend. Edges are labeled with the *cutting cost*, which models the communication cost required to reconcile the state of the blocks across compute nodes in the cluster. The latter is linear to the operation’s inbound dependent data. We assume that

¹The parser accepts unannotated declarations, which are interpreted as `pub0`.

cutting and computation costs are directly comparable through some scaling factor (α), but otherwise omit that detail as it is irrelevant to the function of our compiler.

Lian next uses combinatorial optimization to find the efficient cuts to partition G . This optimization problem is most easily understood as an ILP, though as the constituent variables are all Booleans it is amenable to representation as a PBO problem as well.

We formally give the `opt(·)` function for this problem, as well as all of its constraints, in Sec. 5.2. Informally, the objective can be understood as the maximum of the cumulative computation costs of each partition, plus the sum of all cutting costs on edges bisected by the $\kappa - 1$ cuts (needed to create the user-specified κ chunks). Users can adjust the parameters of the ILP/PBO instances or objective functions to align with different ZK backends and relevant metrics.

The partitioned program is compiled into binaries for the chosen ZKP protocol backend. As Lian is designed to use commit-and-prove ZK protocols as a black box without any internal changes of them, our highly adaptable framework can support a wide range of backends, including the NIZK protocols obtained via the Fiat-Shamir transform [BR93]. These binaries are then sent to any and all individual provers and verifiers. All computation up to this point is based only on public K_0 information, and so the binaries may be freely distributed and reused. However, for the programs to run in parallel, the dependent values between partitions must be computed ahead of time; this is deep simulation. The prover uses their private information to simulate execution of each substatement program, and saves the output values of each one that then become the inputs to others. These output values are the ‘extended witnesses’ of Figure 2. For example, in Figure 1, a cut at L12 would make each entry in \mathbb{M} an extended witness. After deep simulation completes, the prover can, with any verifier, execute the ZKP protocol by running their binaries in parallel and connecting each binary to the other party’s running copy over the network (*i.e.*, `part0` to `part0`, `part1` to `part1`, and so on up to `part κ`).

Finally, after all the chunks have been checked in parallel, their inputs and outputs must be reconciled to check consistency of variables passed across the cuts. This step is essential, otherwise the prover could violate soundness by inconsistently computing their extended witnesses, so that, *e.g.*, $t = 0$ in one partition but $t = 1$ in another. Therefore, the runtime of the proving is not only a function of how large each of the partitioned substatements are, but also of how much data must be communicated amongst the compute cores during this consistency checking step, thus justifying its inclusion in the optimization objective. Lian automatically compiles this reconciliation step into the binaries themselves, so it requires no additional expertise on behalf of the verifier.

3 Ou Language

3.1 Language Syntax

Figure 4 demonstrates Ou’s syntax. Note that we use a^+ as shorthand for one or more non-terminals separated by commas like a_1, \dots, a_n , and use a^* to represent zero or more non-terminals likewise. A program consists of a collection of struct declaration (*sd*), function declaration (*fd*), and external function declaration (*fx*) statements. There must exist a `main` function acting as the entry point to the entire program. The syntax is very similar to many imperative languages, except for the addition of *security levels* ℓ to label primitive types and *annotations* η to label certain functions.

Each atomic type τ is annotated with a security level ℓ . Arrays have both a τ indicating their elements’ type, as well as an ℓ indicating their *access level*. The access level can be viewed as the security level of indexes. A struct’s type only mentions its name, and not its fields, so that we can support recursive struct definitions.

There are two different kinds of expressions: L-expressions (ϵ) and R-expressions (e). An L-expression refers to a memory location like a variable x , an array index $\epsilon[e]$, a struct field $\epsilon.l$, or a dereferenced pointer $*e$. An R-expression represents a value like an integer, float, or Boolean constant; the calling of a builtin


```

 $\ell$  ::= pub0 | pub1 | pub2 | pvt1 | pvt2 | plc1 | plc2
 $\tau$  ::=  $\ell$  int |  $\ell$  bool |  $\ell$  float | unit |  $\tau$  [ $\ell$  num] | struct  $s$  |  $\tau$  *
 $e$  ::=  $\ell$  num |  $\ell$  true |  $\ell$  false | null | tt | ( $e$ )
      |  $f(e^*)$  |  $\ell \{e^+\}$  |  $\{(\text{label} = e)^+\}$  |  $\epsilon$  |  $\& \epsilon$ 
 $\epsilon$  ::=  $x$  |  $\epsilon[e]$  |  $\epsilon.\text{label}$  | *  $e$ 
 $c$  ::=  $\tau x = e$  |  $\epsilon = e$ 
      |  $\tau x = f(e^*)$  |  $c_1; c_2$  |  $\{c\}$  | assert  $e$ 
      | if  $e$  then  $c_1$  else  $c_2$  | while  $e$  do  $c$ 
      | return  $e$  | return | break | continue

 $\eta$  ::= normal | atomic | box1 | box2 | plocal1 | plocal2
 $sd$  ::= struct  $s = \{(\tau \text{label})^+\}$ 
 $fd$  ::=  $\eta \tau f((\tau x)^*)\{c\}$ 
 $fx$  ::=  $\eta \tau f((\tau)^*)$ 
 $prog$  ::=  $(sd|fd|fx)^*$ 

```

Figure 4: Syntax of the Ou language

function, array constructor, or struct constructor; or the loading of an L-expression or a pointer to an L-expression. When constructing an array $\ell \{e_0, \dots, e_{n-1}\}$, the user has to specify its access level. The user can also write normal binary and unary operators, including typecasting following the arrows in Figure 3, and all typical arithmetic operations on integers, floats, and Booleans, but Lian will automatically convert them into builtin function calls. For example, given a variable x of type **pvt1 int** and a variable y of type **pub0 int**, $x + y$ will be automatically rewritten into the expression `pvt1_int_add(x, pvt1_int_of_pub0(y))`.

A command c can be a variable declaration, an assignment, the calling of a user-defined function, conditional branching, a `while` loop, a sequence of commands $c_1; c_2$, a block of commands, `return`-ing from a function, or `break`-ing out of a loop.

Note the syntax treats builtin functions and user-defined functions differently. A builtin function f can be called in expressions freely, while a user-defined function f has to be received immediately by a variable. This distinction is needed because all builtin functions are pure while user-defined functions can have side-effects. As shallow simulation will reorder computations, having side-effects in expressions can alter behavior. We do still allow calling user-defined functions freely in any expression, but just like in C the programmer cannot assume any execution order as Lian will automatically rewrite the code to extract these calls.

Each function also has an annotation η which enforces restrictions on its body. Most common are **normal** functions (the parser interprets unannotated functions as **normal**) which only permit branching on **pub0** Boolean expressions. An **atomic** function is similar, except that shallow simulation will not unfold its body. To branch on a **pub1** value, the user must write the code inside a **box1** function, while **box2** is similar except for additionally allowing branching on **pub2** variables. Functions annotated by **plocal1** or **plocal2** are prover-local functions, which only manipulate **plc1** or **plc2** values respectively. We call functions annotated by **box** or **plocal** *sandboxed* functions, as the type system enforces restrictions to isolate their effects. We further elaborate in Sec. 3.2.

Before typechecking, Lian first scans the program and collects some top-level definitions: Θ maps a struct name s to its definition $\{l_1 : \tau_1, l_2 : \tau_2, \dots\}$. Λ^u maps a user-defined function name f to its definition or declaration. When f is defined, $\Lambda^u(f)$ is like **internal** $(\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c)$ consisting of its annotation, parameter list, return type, and function body; when f is external, $\Lambda^u(f)$ records only **external** $(\eta, (\tau_1, \dots, \tau_n) \rightarrow \tau)$. Note the language does not support global variables, so the prover can only use external functions to communicate her secret with the program. In addition, the language has builtin functions' typing information Λ^b , which map a builtin function name f to its signature $(\tau_1, \dots, \tau_n) \rightarrow \tau$.

Builtins do not have an explicit function body c , but in Sec. 3.3 we will assign each a math interpretation (ff).

In order to ensure the security of ZKP programs, it is crucial for users to label witness variables as `pvt` accurately. Failing to do so can lead to information leakage. `Lian` offers provably-sound typing rules that enable users to implement secure `Ou` programs. By implementing a type checker for these rules, `Lian` can identify and report any errors resulting from incorrect usage of security levels or annotations that violate the typing rules.

3.2 Typing Rules

Figure 5 demonstrates the language’s typing rules. Each typing rule relies on the defined struct definitions Θ and function definitions Λ^u and Λ^b . As these are all fixed we omit them from our written rules for concision. We also rely on a typing environment Γ , which is a list of scopes $\gamma_1 \cdots \gamma_n$. Each scope γ_i maps variable names to their types. All newly defined variables reside in the first scope γ_1 . Note that though the compiler allows the same name to be defined in different scopes, for clarity our discussion assumes that names never conflict.

Typing orders. We say $\ell_1 \leq \ell_2$ if there’s a path from ℓ_1 to ℓ_2 in Figure 3. We say a type τ is *wellformed* if:

- it is an atomic type like ℓ `int` or `bool`;
- it is `struct` s such that s is defined in Θ and the types of all its fields are wellformed; or
- it is an array of type τ [ℓ n] such that τ is wellformed and ℓ is no greater than any security level in τ .

For the last condition the relevant security levels include those of the atomic type as well as the array’s access levels. Wellformedness guarantees the values loaded from arrays are never less secure than the access needed to load them. This restriction is important for both shallow simulation and deep simulation as they need to fully evaluate all K_0 expressions and K_1 expressions, respectively. For example, if $a[e]$ loads a `pub0` integer from array a with `pvt1` access, then shallow simulation does not have enough knowledge to evaluate e , thus it can not know which value is loaded. So although $a[e]$ ’s type is `pub0`, the shallow simulation cannot evaluate it. To make shallow and deep simulation behave as expected, we assume all types in this paper are wellformed.

Typing expressions. For L-expressions, $\Gamma \vdash_L e : \tau, \ell$ means in the typing environment Γ an L-expression e points to a value of type τ , and the L-expression itself has access level ℓ . The access level can help us determine whether an L-expression can be fully evaluated during shallow simulation (`pub0`) or deep simulation (`pub1`, `pvt1` and `plc1`) or neither (`pub2`, `pvt2` and `plc2`). Similarly, for R-expressions, $\Gamma \vdash_R e : \tau$ means an R-expression e is of type τ in the typing environment Γ .

Note that in practice the `Lian` compiler is more forgiving than the language’s strict typing rules, as it performs implicit type conversion along the security lattice’s arrows. For example, an expression e of type `pub0 int` can be used as a `pvt1 int`, and `Lian` will automatically wrap it inside a casting function `pvt1_int_of_pub0(e)`.

We also require all pointers to have `pub0` access level. This restriction is important to the live variable analysis needed for partitioning. Suppose we were to allow a pointer p to have a `pvt1` access level: then shallow simulation could not evaluate the pointer itself in order to determine which location is mutated after running $*p = e$. This would mean all locations in the stack could potentially be the address. Similarly, loading a value from $*p$ would mean any location could potentially be read. Therefore this restriction on pointer access levels is needed to make the live variable analysis yield meaningful results, rather than just concluding by labeling all variables as live.

Typing commands. When typechecking a command c which constitutes a function’s body, we need its return type τ_r to guarantee that each `return` command in c yields a valid value. We also need a set of

$$\boxed{\Gamma \vdash_L \epsilon : \tau, \ell}$$

$$\frac{\frac{x : \tau \in \gamma}{\Gamma \vdash_R e : \tau *}}{\Gamma \vdash_L * e : \tau, \text{pub0}} \quad \frac{\Gamma \vdash_L \epsilon : \tau[\ell n], \ell' \quad \Gamma \vdash_R e : \ell \text{ int} \quad \ell \leq \ell'}{\Gamma \vdash_L \epsilon : \mathbf{struct} s, \ell \quad s : m \in \Theta \quad l : \tau \in m} \quad \frac{\Gamma \vdash_L \epsilon[e] : \tau, \ell}{\Gamma \vdash_L \epsilon.l : \tau, \ell}$$

$$\boxed{\Gamma \vdash_R e : \tau}$$

$$\frac{}{\Gamma \vdash_R \ell n : \ell \text{ int}} \quad \frac{\Gamma \vdash_L \epsilon : \tau, \ell \quad \Gamma \vdash_L \epsilon : \tau, \text{pub0}}{\Gamma \vdash_R \& \epsilon : \tau *}$$

$$\frac{\forall i \in [0, n-1], \Gamma \vdash_R e_i : \tau}{\Gamma \vdash_R \ell \{e_0, \dots, e_{n-1}\} : \tau[\ell n]} \quad \frac{\mathbf{f} \mapsto ((\tau_1, \dots, \tau_n) \rightarrow \tau) \in \Lambda^b \quad \forall i \in [1, n], \Gamma \vdash_R e_i : \tau_i}{\Gamma \vdash_R \mathbf{f}(e_1, \dots, e_n) : \tau}$$

$$\boxed{\Gamma, \tau_r, L_b \vdash c : \Gamma'}$$

$$\frac{\gamma \cdot \Gamma \vdash_R e : \tau}{\gamma \cdot \Gamma, \tau_r, L_b \vdash \tau x = e : \gamma \cup \{x : \tau\} \cdot \Gamma} \quad \frac{f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \text{ or } f \mapsto \mathbf{external} (\eta, (\tau_1, \dots, \tau_n) \rightarrow \tau) \in \Lambda^u \quad \forall i \in [1, n], \gamma \cdot \Gamma \vdash_R e_i : \tau_i}{\gamma \cdot \Gamma, \tau_r, L_b \vdash \tau x = f(e_1, \dots, e_n) : \gamma \cup \{x : \tau\} \cdot \Gamma}$$

$$\frac{\Gamma \vdash_R e : \ell \mathbf{bool}}{\Gamma, \tau_r, L_b \vdash \mathbf{assert} e : \Gamma} \quad \frac{\Gamma \vdash_R e : \ell \mathbf{bool} \quad \Gamma, \tau_r, L_b \vdash c_1 : \Gamma_1 \quad \ell \in L_b \quad \Gamma, \tau_r, L_b \vdash c_2 : \Gamma_2}{\Gamma, \tau_r, L_b \vdash \mathbf{if} e \text{ then } c_1 \text{ else } c_2 : \Gamma}$$

$$\frac{\Gamma \vdash_R e : \tau_r}{\Gamma, \tau_r, L_b \vdash \mathbf{return} e : \Gamma}$$

$$\frac{\Gamma, \tau_r, L_b \vdash c_1 : \Gamma_1 \quad \Gamma_1, \tau_r, L_b \vdash c_2 : \Gamma_2}{\Gamma, \tau_r, L_b \vdash c_1; c_2 : \Gamma_2} \quad \frac{\Gamma \vdash_R e : \ell \mathbf{bool} \quad \ell \in L_b \quad \Gamma, \tau_r, L_b \vdash c : \Gamma'}{\Gamma, \tau_r, L_b \vdash \mathbf{while} e \text{ do } c : \Gamma}$$

Figure 5: Typing Rules

branching security levels L_b to guard all branching expressions inside c . In **normal** and **atomic** functions, all branching expressions have to be **pub0**, but prover-local and boxed functions enforce fewer restrictions. **Typing a program.** The type checker first checks all struct definitions, and makes sure there are no name collisions between struct names and field names, and then checks each function’s body. For a function definition $\eta \tau f(\tau_1 x_1, \dots, \tau_n x_n)\{c\}$, we apply the command’s typing rule on its body: $\Gamma, \tau, L_b \vdash \{c\} : \Gamma$, where $\Gamma = \{x_1 : \tau_1, \dots, x_n : \tau_n\}$ and L_b is determined by η :

$$L_b = \begin{cases} \{\text{pub0}\} & \text{if } \eta \text{ is } \mathbf{atomic} \text{ or } \mathbf{normal} \\ \{\text{pub0}, \text{pub1}\} & \text{if } \eta \text{ is } \mathbf{box1} \\ \{\text{pub0}, \text{pub1}, \text{pub2}\} & \text{if } \eta \text{ is } \mathbf{box2} \\ \{\text{plc1}\} & \text{if } \eta \text{ is } \mathbf{plocal1} \\ \{\text{plc1}, \text{plc2}\} & \text{if } \eta \text{ is } \mathbf{plocal2}. \end{cases}$$

There are additional restrictions placed on sandboxed functions. We require that **plocal1** and **box1** functions only modify external data that is K_1 or K_2 , and only return K_1 or K_2 results. Similarly **plocal2** and **box2** functions can only modify external data that is K_2 , and only produce K_2 results. The reason is similar to the restrictions on arrays’ access levels: shallow simulation cannot execute **plocal1** or **box1** functions, so it is infeasible to allow them to modify or return any K_0 data. If they could, simulation of any other computations that rely on such data could not proceed after the function call. We also restrict that **plocal** functions can only manipulate **plc** data. Since all **plocal** functions only run on the prover’s side they cannot include any **pub** or **pvt** values, as their use would require the verifier’s participation.

To enforce the first restriction, we only need to examine the function’s signature: all pointers passed into **plocal1** and **box1** functions should only point to K_1 or K_2 data, while all pointers passed into **plocal2** and **box2** functions should only point to K_2 data. The return types have similar restrictions. Since the language does not support global variables, sandboxed functions can only update external data via the pointers in the arguments. To enforce the second restriction, we need to examine the entire function body to make sure all expressions’ types only have **plc** annotations, and the body only calls other **plocal** functions.

3.3 Dynamic Semantics

This section defines the dynamic semantics that apply if all secrets are known. The semantics correspond to the program’s behavior in the ideal world where it has enough knowledge to fully evaluate all expressions into values and check if all assertions are satisfied. We will use this semantics as a baseline to prove the simulation and distribution are correct. The semantics rules are not surprising. The detailed definitions are listed in appendix B.

R-values and L-values. We use v to denote R-values (or “values” for short) and μ to denote L-values. R-values can be atomic values like **Vunit** and **Vint** ℓn , an empty pointer **Vnull**, a pointer **Vref** μ , an array **Varray** ℓa where ℓ is the array’s access level and a is a mapping from indexes to R-values like $\{0 \mapsto v_0, \dots, n \mapsto v_n\}$, or a struct **Vstruct** m where m is a mapping from label to R-values like $\{l_0 \mapsto v_0, \dots, l_n \mapsto v_n\}$. An L-value is a reference to a memory location: it can be **Vvar** α that points to an *alias*, or **Vindex** $\mu \ell i$ that refers to an array μ ’s index i , or **Vfield** $\mu \ell$ that refers to a struct μ ’s field ℓ . An alias is used by the stack to disambiguate during recursion, and can be viewed as a memory address.

$$\begin{aligned} v ::= & \mathbf{Vunit} \mid \mathbf{Vnull} \mid \mathbf{Vref} \mu \\ & \mid \mathbf{Vint} \ell \text{ num} \mid \mathbf{Vfloat} \ell \text{ num} \mid \mathbf{Vbool} \ell \text{ bool} \\ & \mid \mathbf{Varray} \ell \text{ array} \mid \mathbf{Vstruct} \text{ map} \\ \mu ::= & \mathbf{Vvar} \alpha \mid \mathbf{Vindex} \mu \ell \text{ num} \mid \mathbf{Vfield} \mu \text{ label} \end{aligned}$$

The program runs on a stack Ω , *i.e.*, a list of frames $\omega_1 \cdot \omega_2 \cdots \omega_n$. The last frame ω_n stores **main**’s variables. The stack grows with nested function calls, and the top frame ω_1 stores the current function call’s

local variables. Each frame ω is a map from aliases to values. We use $\Omega[\alpha \mapsto v]$ to denote adding a mapping $\alpha \mapsto v$ to the top frame. Whenever a new variable x is defined, `Lian` allocates a unique alias $\text{new}(x)$ for it. This is useful when recursion occurs, as the same variable x might have different instances in the stack, so we cannot bind their values to the same name x . We also use a function $\text{find}(\Omega, x)$ to find the latest alias allocated for x in the stack. Note that when x is defined for the first time, $\text{new}(x)$ returns x directly instead of allocating a new alias for it.

Given a stack Ω , we use a function $\text{load}(\Omega, \mu)$ to load the value referred by μ in Ω , and use a function $\text{store}(\Omega, \mu, v)$ to compute the updated stack after replacing the value at μ with v . They both search the entire stack from top frame to bottom instead of working only with the top frame.

Semantics. The semantics look up the sets of user-defined functions Λ^u and builtin functions Λ^b to interpret function calls. As they are immutable, for concision we omit them from our rules.

Expressions: $\llbracket \epsilon \rrbracket_L^\Omega$ computes L-expression ϵ in stack Ω and yields an L-value, $\llbracket e \rrbracket_R^\Omega$ evaluates R-expression e and yields an R-value. Note `Lian` assigns a mathematical interpretation $\llbracket f \rrbracket$ for each builtin function f which maps a tuple of R-values to an R-value. For example, $\llbracket \text{pvt1_int_add} \rrbracket(\mathbf{Vint\ pvt1\ } n, \mathbf{Vint\ pvt1\ } m) = \mathbf{Vint\ pvt1\ } (n + m)$ is assigned by the compiler for integer addition.

Commands: $(c, \Omega_1) \rightarrow_C (r, \Omega_2, \beta)$ means that running a command c from a stack Ω_1 yields a sequence of assertion results β , a result r which is either `cont` or `ret v`, and a changed stack Ω_2 . Since the command constitutes a function's body, we use `ret v` to mean the function returns with v and `cont` to mean it has not yet returned.

4 Shallow Simulation

Shallow simulation executes and unfolds the program with K_0 knowledge to obtain a simplified sequential program. `Lian` will attempt as much simplification as possible, and output all unknown computations as a sequence of commands for later analysis and partitioning. All **normal** and **atomic** functions will be executed, but only the former's execution history will be output, while the latter will be output as indivisible calls. As sandboxed functions rely on K_1 or K_2 knowledge, they will not be executed in shallow simulation and will be kept in the generated sequential program.

4.1 Shallow Semantics

This section formalizes the shallow simulation's behavior as a *shallow semantics*. We use symbolic values to define its semantics.

$$\begin{aligned}
 v ::= & \mathbf{SVunit} \mid \mathbf{SVnull} \mid \mathbf{SVref\ } \mu \\
 & \mid \mathbf{SVint\ } \ell \text{ num} \mid \mathbf{SVbool\ } \ell \text{ bool} \\
 & \mid \mathbf{SVfloat\ } \ell \text{ num} \mid \mathbf{SVsym} \\
 & \mid \mathbf{SVarray\ } \ell \text{ array} \mid \mathbf{SVstruct\ } \text{map}
 \end{aligned}$$

Figure 6 shows its details. The actual implementation uses several optimizations to speed up the simulation.

Symbolic values. To define the semantics of both shallow simulation and deep simulation, we define *symbolic values* to represent R-values not known during the simulation. We do so by redefining the R-value constructor from Sec. 3.3 to support **SVsym** variants, and also use *symbolic stacks*, which store symbolic values.

Builtin functions. During shallow simulation only a subset of the builtin functions can be interpreted as the rest rely on either K_1 or K_2 knowledge. We use $\llbracket f \rrbracket_0$ to denote a shallow simulatable interpretation of a builtin function. It takes in a list of expressions and yields an expression as the result. For example, we have the interpretation $\llbracket \text{pub0_int_add} \rrbracket_0(\text{pub0\ } 1, \text{pub0\ } 2) = \text{pub0\ } 3$, while $\text{pvt1_int_add} \notin \llbracket - \rrbracket_0$ as it relies on K_1 knowledge.

Semantics. $\llbracket \epsilon \rrbracket_{L_0}^\Omega = \epsilon'$ means simplifying an L-expression ϵ in a symbolic stack Ω with K_0 knowledge results in an L-expression ϵ' . These semantics simplify an L-expression instead of evaluating it directly,

as the stack may not contain enough information to fully evaluate it to an L-value. Similarly, $\llbracket e \rrbracket_{R_0}^\Omega = e'$ simplifies an R-expressions.

When storing a simplified expression into the stack, we use a partial function $\llbracket - \rrbracket_{L_0}$ to evaluate an L-expression down to an L-value, and a total function $\llbracket - \rrbracket_{R_0}$ to evaluate an R-expression down to a symbolic R-value. $\llbracket - \rrbracket_{R_0}$ is able to be total because it replaces all unknown subexpressions by **SVsym**. When loading a value from the stack, we use a total function $\lceil - \rceil_L$ to lift a symbolic L-value to an L-expression, and a total function $\lceil -, - \rceil_R$ to lift a symbolic R-value to an R-expression and concretize any **SVsym** variables using the appropriate L-value. See Figure 19.

We reuse `load` and `store` to load and store symbolic values to and from the symbolic stack.

$$\text{sload}(\Omega, \epsilon) = \begin{cases} \lceil \text{load}(\Omega, \mu), \mu \rceil_R & \text{when } \llbracket \epsilon \rrbracket_{L_0} = \mu \\ \epsilon & \text{otherwise} \end{cases}$$

$$\text{sstore}(\Omega, \epsilon, e) = \begin{cases} \text{store}(\Omega, \mu, \llbracket e \rrbracket_{R_0}) & \text{when } \llbracket \epsilon \rrbracket_{L_0} = \mu \\ \Omega & \text{otherwise} \end{cases}$$

When ϵ is not completely known, type wellformedness guarantees that the symbolic value referred to by it must also be unknown, *i.e.*, **SVsym** or an array or struct of unknown values. When so the value to be stored must also be unknown, thus `sstore` behaves correctly even when ϵ is not fully evaluated.

Commands: $(c, \Omega_1) \rightsquigarrow_C (r, \Omega_2, h)$ means simplifying a command c in a symbolic stack Ω_1 results in: a return status r which is either `cont` or `ret e`, a new stack Ω_2 , and a sequence of history h . We use \cdot to denote an empty history.

Shallow simulation does not have enough knowledge to simulate a sandboxed function, thus it returns a symbolic value `default(τ)` generated from the type τ . Also, the typing rule guarantees even if a pointer is passed into the function, the pointed data must be filled with **SVsym**. So no matter whether the function has side effect or not, the symbolic stack remains unchanged.

Note that we have added two pseudo-commands **push** and **pop** into the generated program. They help mark the points where a function is called and returned, and their dynamic semantics are just pushing and popping stacks. Though helpful for the compiler, ignoring them does not change the program's behavior as the alias mechanism guarantees there will never be name conflicts. We introduces them mainly as machinery to help us prove refinement.

Atomic function calls. The size of a generated program can be very large as it is linear in the execution time. To reduce the challenge for partitioning, we allow users to annotate some functions as **atomic** in order to contract the function call into one command in the generated program. Atomic annotations are optional for programs and are intended solely for optimization purposes. The effectiveness of an atomic annotation depends on the user's understanding of the program structure. If an atomic annotation is inappropriate, it may increase the compilation time, but it will *not* compromise the program's security.

Their semantics are very similar to **normal** function calls, but instead of generating a sequence of history, an atomic function call $\tau x = f(e_1, \dots, e_n)$ only yields one command in the history: $\tau \alpha = f(e'_1, \dots, e'_n) \mathbf{R} S_R \mathbf{W} S_W$. Here we add a new kind of command to encode atomic function calls, but its dynamic semantics are the same as for normal function calls. The two extra sets S_R and S_W are only used to help perform live variable analysis.

S_R records a set of L-values defined before calling f that may be read during the function call, while S_W records a set of L-values defined before calling f that may be updated during the function call. They are collected during the simulation using functions from Sec. 5.1: $\text{REF}(c)$ computes all L-values that may be read in command c ; $\text{DEF}(c)$ computes all L-values that may be updated in c .

One convenient fact about shallow simulation is that no pointer dereference will appear in a generated program. It is a corollary of a more essential property: shallow simulation fully evaluates all expressions

$$\begin{array}{l}
\llbracket x \rrbracket_{L_0}^\Omega = \text{find}(\Omega, x) \\
\llbracket \epsilon[e] \rrbracket_{L_0}^\Omega = \llbracket \epsilon \rrbracket_{L_0}^\Omega \llbracket [e] \rrbracket_{R_0}^\Omega \\
\llbracket \epsilon.l \rrbracket_{L_0}^\Omega = \llbracket \epsilon \rrbracket_{L_0}^\Omega.l \\
\llbracket * e \rrbracket_{L_0}^\Omega = \epsilon \quad \text{where } \llbracket e \rrbracket_{R_0}^\Omega = \& \epsilon \\
\llbracket \ell n \rrbracket_{R_0}^\Omega = \ell n \\
\llbracket \epsilon \rrbracket_{R_0}^\Omega = \text{sload}(\Omega, \llbracket \epsilon \rrbracket_{L_0}^\Omega) \\
\llbracket \& \epsilon \rrbracket_{R_0}^\Omega = \& \llbracket \epsilon \rrbracket_{L_0}^\Omega \\
\llbracket f(e_1, \dots, e_n) \rrbracket_{R_0}^\Omega = \begin{cases} (\text{f})_0(\llbracket e_1 \rrbracket_{R_0}^\Omega, \dots, \llbracket e_n \rrbracket_{R_0}^\Omega) \\ \text{when } \text{f} \in \text{dom}(\text{--})_0 \\ f(\llbracket e_1 \rrbracket_{R_0}^\Omega, \dots, \llbracket e_n \rrbracket_{R_0}^\Omega) \quad \text{o/w} \end{cases} \\
\llbracket \mathbf{Svint} \ell n, \mu \rrbracket_R = \ell n \\
\llbracket \mathbf{SVref} \mu', \mu \rrbracket_R = \& \llbracket \mu' \rrbracket_L \\
\llbracket \mathbf{SVsym}, \mu \rrbracket_R = \llbracket \mu \rrbracket_L \\
\llbracket \mathbf{SVarray} \ell a, \mu \rrbracket_R = \ell \{ \llbracket a(0), \mathbf{Vindex} \mu \ell 0 \rrbracket_R, \dots \} \\
\llbracket \mathbf{SVstruct} m, \mu \rrbracket_R = \{ .l_1 = \llbracket m(l_1), \mathbf{Vfield} \mu l_1 \rrbracket_R, \dots \} \\
\llbracket \mathbf{Vvar} \alpha \rrbracket_L = \alpha \\
\llbracket \mathbf{Vindex} \mu \ell n \rrbracket_L = \llbracket \mu \rrbracket_L \llbracket \ell n \rrbracket \\
\llbracket \mathbf{Vfield} \mu l \rrbracket_L = \llbracket \mu \rrbracket_L.l \\
\llbracket \alpha \rrbracket_{L_0} = \mathbf{Vvar} \alpha \\
\llbracket \epsilon[\ell n] \rrbracket_{L_0} = \mathbf{Vindex} \mu \ell n \text{ when } \llbracket \epsilon \rrbracket_{L_0} = \mu \text{ and } \ell = \text{pub0} \\
\llbracket \epsilon.l \rrbracket_{L_0} = \mathbf{Vfield} \mu l \quad \text{when } \llbracket \epsilon \rrbracket_{L_0} = \mu \\
\llbracket * (\& \epsilon) \rrbracket_{L_0} = \mu \quad \text{when } \llbracket \epsilon \rrbracket_{L_0} = \mu \\
\llbracket \ell n \rrbracket_{R_0} = \begin{cases} \mathbf{Svint} \ell n & \text{when } \ell = \text{pub0} \\ \mathbf{SVsym} & \text{o/w} \end{cases} \\
\llbracket \& \epsilon \rrbracket_{R_0} = \mathbf{SVref} \llbracket \epsilon \rrbracket_{L_0} \\
\llbracket \ell \{e_0, \dots\} \rrbracket_{R_0} = \mathbf{SVarray} \ell \cup \{i \mapsto \llbracket e_i \rrbracket_{R_0}\} \\
\llbracket \{.l_1 = e_1, \dots\} \rrbracket_{R_0} = \mathbf{SVstruct} \cup \{l_i \mapsto \llbracket e_i \rrbracket_{R_0}\} \\
\llbracket \epsilon \rrbracket_{R_0} = \mathbf{SVsym} \\
\llbracket f(e_1 \dots e_n) \rrbracket_{R_0} = \mathbf{SVsym} \\
\boxed{(c, \Omega) \rightsquigarrow_C (r, \Omega', h)} \\
\frac{\text{new}(x) = \alpha \quad \llbracket e \rrbracket_{R_0}^\Omega = e'}{(\tau x = e, \Omega) \rightsquigarrow_C (\text{cont}, \Omega[\alpha \mapsto \llbracket e \rrbracket_{R_0}^\Omega], \tau \alpha = e')} \quad \frac{\llbracket \epsilon \rrbracket_{L_0}^\Omega = \epsilon' \quad \llbracket e \rrbracket_{R_0}^\Omega = e'}{(\epsilon = e, \Omega) \rightsquigarrow_C (\text{cont}, \text{sstore}(\Omega, \epsilon', e'), \epsilon' = e')} \\
\frac{}{(\mathbf{return} e, \Omega) \rightsquigarrow_C (\text{ret} \llbracket e \rrbracket_{R_0}^\Omega, \Omega, \cdot)} \quad \frac{}{(\mathbf{assert} e, \Omega) \rightsquigarrow_C (\text{cont}, \Omega, \mathbf{assert} \llbracket e \rrbracket_{R_0}^\Omega)} \\
\frac{\llbracket e \rrbracket_{R_0}^\Omega = \text{pub0} \text{ true} \quad (c_1, \Omega) \rightsquigarrow_C (r, \Omega', h)}{(\mathbf{if} e \text{ then } c_1 \text{ else } c_2, \Omega) \rightsquigarrow_C (r, \Omega', h)} \quad \frac{(c_1, \Omega) \rightsquigarrow_C (\text{cont}, \Omega_1, h_1) \quad (c_2, \Omega_1) \rightsquigarrow_C (r, \Omega_2, h_2)}{(c_1; c_2, \Omega) \rightsquigarrow_C (r, \Omega_2, h_1; h_2)} \\
\frac{\text{new}(x) = \alpha \quad \eta = \mathbf{normal} \\ f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \\ \forall i \in [1, n], \llbracket e_i \rrbracket_{R_0}^\Omega = e'_i \wedge \text{new}(x_i) = \alpha_i \\ (c, \cup \{\alpha_i \mapsto \llbracket e'_i \rrbracket_{R_0}^\Omega\} \cdot \Omega) \rightsquigarrow_C (\text{ret} e, \omega \cdot \Omega', h) \\ h' = (\mathbf{push}; \tau_1 \alpha_1 = e'_1; \dots; \tau_n \alpha_n = e'_n; h; \mathbf{pop} \tau \alpha = e)}{(\tau x = f(e_1, \dots, e_n), \Omega) \rightsquigarrow_C (\text{cont}, \Omega'[\alpha \mapsto \llbracket e \rrbracket_{R_0}^\Omega], h')} \\
\frac{\text{new}(x) = \alpha \quad \eta \notin \{\mathbf{normal}, \mathbf{atomic}\} \\ f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \\ \text{or } f \mapsto \mathbf{external} (\eta, (\tau_1, \dots, \tau_n) \rightarrow \tau) \in \Lambda^u \\ \forall i \in [1, n], \llbracket e_i \rrbracket_{R_0}^\Omega = e'_i \\ \Omega' = \Omega[\alpha \mapsto \text{default}(\tau)]}{(\tau x = f(e_1, \dots, e_n), \Omega) \rightsquigarrow_C (\text{cont}, \Omega', \tau \alpha = f(e'_1, \dots, e'_n))}
\end{array}$$

Figure 6: Shallow Semantics

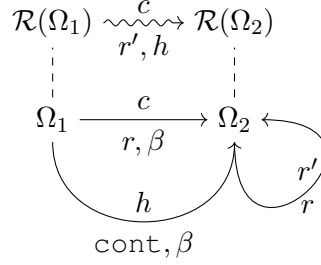


Figure 7: Shallow Semantics Refinement

that only need K_0 knowledge, including all pointers. In Appendix C, we define *normalized* expressions to formally describe this property to help prove the correctness of shallow simulation.

4.2 Correctness of Shallow Simulation

We start by defining a refinement function \mathcal{R} mapping R-values to symbolic R-values. It replaces all non-`pub0` atomic values with **SVsym**. We overload \mathcal{R} to also represent the function that performs the transformations to convert a stack into a symbolic stack.

$$\begin{aligned} \mathcal{R}(\mathbf{Vint} \ell n) &= \begin{cases} \mathbf{SVint} \ell n & \text{when } \ell = \text{pub0} \\ \mathbf{SVsym} & \text{otherwise} \end{cases} \\ \mathcal{R}(\mathbf{Vref} \mu) &= \mathbf{SVref} \mu \\ \mathcal{R}(\mathbf{Varray} \ell a) &= \mathbf{SVarray} \ell \cup \{i \mapsto \mathcal{R}(a(i))\} \\ \mathcal{R}(\mathbf{Vstruct} m) &= \mathbf{SVstruct} \cup \{l_i \mapsto \mathcal{R}(m(l_i))\} \end{aligned}$$

See Figure 7. We also define a semantics $\llbracket r \rrbracket_R^\Omega$ in order to evaluate return statuses.

$$\llbracket \text{cont} \rrbracket_R^\Omega = \text{cont} \qquad \llbracket \text{ret } e \rrbracket_R^\Omega = \text{ret } \llbracket e \rrbracket_R^\Omega$$

Theorem 1 *If $(c, \Omega_1) \rightarrow_C (r, \Omega_2, \beta)$, then there exist r' and h s.t.*

$$\begin{aligned} (c, \mathcal{R}(\Omega_1)) &\rightsquigarrow_C (r', \mathcal{R}(\Omega_2), h) \\ \text{and } (h, \Omega_1) &\rightarrow_C (\text{cont}, \Omega_2, \beta) \text{ and } \llbracket r' \rrbracket_R^{\Omega_2} = r. \end{aligned}$$

Figure 7 demonstrates this theorem. Although shallow simulation reorders some computations by simplifying K_0 expressions and unfolding the program, this theorem guarantees such reordering does not change the behavior. We prove the theorem in Appendix C.2.

5 Statement Partitioning

The shallow simulation unfolds the original program into a sequential program $h = c_1; \dots; c_n$. The compiler then tries to find a way to partition these commands into κ chunks so that the distributed running time and the communication time are minimized. Sec. 5.1 describes the live variable analysis algorithm that finds data dependencies in the sequential program. Sec. 5.2 then describes how **Lian** encodes the cut search problem and uses a solver to find the efficient partition. Sec. 5.3 then explains how we generate the distributed program from the sequential program, the partition, and the dependencies.

5.1 Live Variable Analysis

Lian performs live variable analysis on a sequential program $h = c_1; \dots; c_n$. The goal is to generate a dependency graph $DG = (DV, DE)$ where $DV = \{1 \dots n\}$ and each edge $(i, j) \in DE$ is decorated with a set of L-values $\text{DEP}(i, j)$ that *may* be read in c_j and *may* be manipulated in c_i . We use ‘may’ rather than ‘will’ because shallow simulation cannot determine all L-expressions’ values. Nonetheless, we can

still conduct an over-approximate analysis using some approximation functions. Note $\text{DEP}(i, j)$ is only meaningful when $i < j$, so $\text{DEP}(i, j) = \emptyset$ when $i \geq j$.

Appendix D defines two utility functions: $\text{REF}(c)$ computes the set of L-values that may be read in c , while $\text{DEF}(c)$ computes the set of L-values that are newly defined or may be updated in c . Recall that we used these functions in Sec. 4.1.

Analysis Algorithm. The program is a sequence of commands $h = c_1; \dots; c_n$. We start live variable analysis from its end and let $\text{LIVE}(n) := \emptyset$. We then scan from the end to the beginning and compute unresolved dependent variables for each command c_i .

$$\begin{aligned} \text{LIVE}(i-1) = & \{\mu \mapsto j \in \text{LIVE}(i) \mid \mu \notin \text{DEF}(c_i) \cup \text{REF}(c_i)\} \\ & \cup \{\mu \mapsto i \mid \mu \in \text{REF}(c_i)\} \end{aligned}$$

$\text{LIVE}(i)$ is a mapping from L-values to the commands that may read them in the future. In the first line, if c_i uses an L-value μ that is later dependent by c_j ,² then we can safely delete it from $\text{LIVE}(i)$ and add that L-value as a dependency between c_i and c_j . However, c_i may read other L-values, so in the second line we mark all L-values from $\text{REF}(c_i)$ as unresolved and dependent on c_i . Such dependencies will eventually be resolved because every L-value is part of a variable defined in the program, and the command that defines the variable cannot read it. Now we can construct the graph DG by defining the dependency function:

$$\text{DEP}(i, j) = \{\mu \mid \mu \mapsto j \in \text{LIVE}(i) \wedge \mu \in \text{DEF}(c_i) \cup \text{REF}(c_i)\}$$

In order to capture all dependencies of and from a command, we denote $\text{DEP}(i, -) := \bigcup_{j=1}^n \text{DEP}(i, j)$ and $\text{DEP}(-, j) := \bigcup_{i=1}^n \text{DEP}(i, j)$. It follows from our definition that DG is a directed acyclic graph, i.e., $i \geq j \implies \text{DEP}(i, j) = \emptyset$. Any L-value dependent on c_i must either be read or written in c_i , any L-value dependent by c_j must be added into the live set during the scan, which can only happen when it is read in c_j , so we have the following lemmas:

Lemma 2 $\text{DEP}(i, -) \subseteq \text{DEF}(c_i) \cup \text{REF}(c_i)$.

Lemma 3 $\text{DEP}(-, j) = \text{REF}(c_j)$.

We need to define a notion of similarity between stacks to reason about the correctness of live variable analysis and program distribution. Given two stacks Ω_1 and Ω_2 and a set of L-values S , we say Ω_1 and Ω_2 are *similar* up to S , denoted as $\Omega_1 \approx_S \Omega_2$, if 1. $\forall \mu \in S, \text{load}(\Omega_1, \mu) = \text{load}(\Omega_2, \mu)$, and 2. the two stacks have similar structures in terms of S , i.e., every alias appearing in S resides in the same frame in both stacks.

Theorem 4 shows the live variable analysis is correct. It says the dependent L-values of c_j are enough to safely execute the command and yield all values that are dependent later. It can be proven using lemma 2 and 3, then prove the values in $\text{REF}(c_j)$ are enough for c_j to execute and produce the desired behavior.

Theorem 4 *If both $(c_1; \dots; c_{j-1}, \Omega) \rightarrow_C (\text{cont}, \Omega_{j-1}, \beta_{j-1})$ and $(c_j, \Omega_{j-1}) \rightarrow_C (\text{cont}, \Omega_j, \beta)$, then*

$$\begin{aligned} \forall \Omega'_{j-1}, \Omega_{j-1} \approx_{\text{DEP}(-, j)} \Omega'_{j-1} \implies \\ \exists \Omega'_j, (c_j, \Omega'_{j-1}) \rightarrow_C (\text{cont}, \Omega'_j, \beta) \wedge \Omega_j \approx_{\text{DEP}(j, -)} \Omega'_j. \end{aligned}$$

²In traditional live variable analysis, a live variable is only resolved when it is modified or defined, but here we also resolve it when it is read. This is a compromise to simplify the efficient-cut search and to reduce cutting cost in most cases. See Appendix D.

5.2 Efficient-Cut Search

The final intermediary pass of the compilation partitions the program into κ chunks for parallel execution. Recall that after the flattening during shallow simulation (Sec. 4) and the live variable analysis (Sec. 5.1), Lian obtains a sequential program $h = c_1; \dots; c_n$ and a dependency graph $DG = (DV, DE)$ where each edge (i, j) in DE is decorated with a set of dependent L-values $\text{DEP}(i, j)$. The compiler then extracts a directed acyclic graph $G = (V, E)$ from DG capturing the dependency relation and the costs. Each node i in V is decorated with c_i 's computation cost PC_i . An edge $(i, j) \in E$ indicates c_j reads at least one L-value used last by c_i , and is decorated with the communication cost C_{ij} and a cuttable bit B_{ij} . If there is a K_2 L-value in $\text{DEP}(i, j)$, then c_i and c_j must reside in the same chunk, thus $C_{ij} = \infty$ and $B_{ij} = 0$; Otherwise, C_{ij} equals to the accumulated cutting cost of all K_1 L-values in $\text{DEP}(i, j)$ and $B_{ij} = 1$.

We use combinatorial optimization constrained by G to find the efficient cutting of h into κ chunks $\langle h_1^{\text{sub}}, h_2^{\text{sub}}, \dots, h_\kappa^{\text{sub}} \rangle$. We natively write this optimization problem as an integer linear program (ILP). However, the variables (though not the constants) it uses are all Booleans, and therefore the objective can be interpreted as a pseudo-Boolean function and optimized using pseudo-Boolean optimization (PBO) as well.

We begin by defining $ub = \log(\sum_{i=1}^n PC_i)$, the length of the cost of the sequential computation. We next define a pair of matrices. $\mathbf{Y} \in \{0, 1\}^{n \times n}$, indicates the location of cuts: $\mathbf{Y}_{ij} = 1$ means that a cut is made between the i -th and j -th instruction block. $\mathbf{X} \in \{0, 1\}^{n \times \kappa}$ indicates the inclusion of instruction blocks into partitions: $\mathbf{X}_{it} = 1$ indicates the i -th instruction is in the t -th partition for $t \in [\kappa]$.

Finding the cheapest partitioning then reduces to solving the optimization objective of

$$\text{obj}(G) = \min_{\mathbf{X}, \mathbf{Y}} \left[\sum_{(i,j) \in E} \mathbf{Y}_{ij} \cdot C_{ij} + \max_{t \in [\kappa]} \sum_{i=1}^n \mathbf{X}_{it} \cdot PC_i \right]$$

subject to some consistency constraints on \mathbf{X} and \mathbf{Y} . In order to soundly model this problem for constrained optimization without a min-max objective requires, for PBO in particular, a bit of due care. To convert the original objective into a PBO problem, here we add a bit-vector \mathbf{COP} to encode $\max_{t \in [\kappa]} \sum_{i=1}^n \mathbf{X}_{it} \cdot PC_i$. Our approach is closely related to prior work in the literature on PBO for graph partitioning [CR93]:

$$\begin{aligned} \min \quad & \sum_{i=1}^{ub} 2^{(i-1)} \cdot \mathbf{COP}_i + \sum_{e \in E} C_e \cdot \mathbf{Y}_e \\ \text{s.t.} \quad & \forall i \in [n] : \quad \sum_{t=1}^{\kappa} \mathbf{X}_{it} = 1 \\ & \forall (i, j) \in E : \quad \mathbf{Y}_{ij} \leq B_{ij} \\ & \forall (i, j) \in E, t \in [\kappa] : \quad \mathbf{X}_{it} + \mathbf{X}_{jt} \leq 2 - \mathbf{Y}_{ij} \\ & \forall (i, j) \in E, t \in [\kappa] : \quad |\mathbf{X}_{it} - \mathbf{X}_{jt}| \leq \mathbf{Y}_{ij} \\ & \forall t \in [\kappa] : \quad \sum_{\ell=1}^{ub} 2^{(\ell-1)} \cdot \mathbf{COP}_\ell \geq \sum_{i=1}^n \mathbf{X}_{it} \cdot PC_i \\ & \forall i \in [n], t \in [\kappa] : \quad \mathbf{X}_{it} \in \{0, 1\} \\ & \forall (i, j) \in E : \quad \mathbf{Y}_{ij} \in \{0, 1\} \\ & \forall \ell \in [ub] : \quad \mathbf{COP}_\ell \in \{0, 1\} \end{aligned}$$

The first constraint enforces that every instruction appears in exactly one partition, the second that only cuttable edges are cut, and the third and fourth together that if the i -th and j -th instruction blocks are in different partitions, then any edge (i, j) must be represented in the appropriate cut by \mathbf{Y}_{ij} . The final four constraints simply enforce the Boolean character of the variables.

The fifth constraint is the most nuanced. It enforces that

$$\text{COP}_i = \text{the } i\text{-th bit of } \max_{t \in [\kappa]} \sum_{i=1}^n \mathbf{X}_{it} \cdot PC_i$$

so that the computation in the objective function is correct. This relationship can be encoded much more simply in a general ILP, but is necessary for PBO compatibility. Notice that no constraints force the partitions to be contiguous: even though h is a sequential program, the h_t^{sub} can have a much more complex structure, at the cost of additional extended witnesses and so more communication during the consistency checking.

With this encoding, `Lian` hands this optimization problem to either a PBO or ILP solver (we use the Gurobi PBO solver in our benchmarks). The solver’s output may not always be optimal in practice, as it is limited by the available computational resources. Consequently, the corresponding cut scheme we employ may also be suboptimal. The compiler then converts the result into a function `chunk(-)` such that `chunk(i) = t` if $X_{it} = 1$ in the solver’s result.

Fact 1 $\text{chunk}(i) \neq \text{chunk}(j) \implies \text{no } K_2 \text{ L-value in } \text{DEP}(i, j)$.

This follows directly from the PBO constraints.

5.3 Distributed Program

Given a sequential program $h = c_1; \dots; c_n$ from the shallow simulation, a dependency graph $\text{DEP}(-, -)$ from live variable analysis, and a partition `chunk(-)` from efficient-cut search, the compiler generates κ programs for distributed execution among κ pairs of prover and verifier computation cores.

The t -th chunk is a sequential program `synct1; ct1; synct2; ct2; ...` such that $t = \text{chunk}(t_1) = \text{chunk}(t_2) = \dots$ and $t_1 < t_2 < \dots$. We denote a communication between the prover and the verifier by `synctj`. These communication rounds are used to share K_0 or authenticate K_1 data that are needed by c_{t_j} , *i.e.*, data pointed by L-values in $\bigcup_{\text{chunk}(i) \neq t} \text{DEP}(i, t_j)$, as necessary. Deep simulation (Sec. 6) computes all this dependent data and distribute it to all of the prover cores. So each pair of prover core and verifier core can run their chunk without the others’ involvement. At the end, all verifier cores will work together to do a consistency check. The goal is to make sure all data provided by the prover cores in `synctj` are indeed the same as produced by their dependent commands in other chunks. Note the deep simulation only computes K_0 and K_1 values, but Fact 1 guarantees all data needed for multiple chunks does not contain K_2 information, thus the deep simulation is enough.

6 Deep Simulation

The deep simulation executes the program with K_0 and K_1 knowledge and evaluates all K_0 and K_1 values. The goal is to compute all data that are dependent between different chunks in the distributed program. It is similar to shallow simulation, except that (i) the definitions of all external functions are known, (ii) the knowledge level is K_1 instead of K_0 , and (iii) the purpose is to evaluate all K_0 and all K_1 values instead of unfolding the program.

6.1 Deep Semantics

This section formalizes deep simulation as a *deep semantics*. We use the same symbolic values and stacks as in Sec. 4.1. We still use $\lceil -, - \rceil_R$ and $\lceil -, - \rceil_L$ to lift R-values and L-values to R-expressions and L-expressions. But when converting in the reversed direction, we use $\lfloor - \rfloor_{R_1}$ and $\lfloor - \rfloor_{L_1}$ instead of $\lfloor - \rfloor_{R_0}$ and $\lfloor - \rfloor_{L_0}$. The only difference is when converting atomic expressions. For example, $\lfloor \ell \ n \rfloor_{R_1} = \mathbf{svint} \ \ell \ n$ when $\ell \leq \text{plc1}$. We also define `dload` and `dstore` based on these operations just like `sload` and `sstore`.

The semantics for L-expressions $\llbracket \epsilon \rrbracket_{L_1}^\Omega$ and R-expressions $\llbracket e \rrbracket_{R_1}^\Omega$ are also similar except for using $\lfloor - \rfloor_{L_1}$ and $\lfloor - \rfloor_{R_1}$ to interact with the symbolic stack, allowing dereference expression to be not fully evaluated,

and use $\llbracket \cdot \rrbracket_1$ to interpret builtin functions. $\llbracket - \rrbracket_1$ is a superset of $\llbracket - \rrbracket_0$ by adding more functions that can be computed with K_1 knowledge. For example, $\llbracket \text{pvt1_int_add} \rrbracket_1(\text{pvt1 } n, \text{pvt1 } m) = \text{pvt1 } (n + m)$. Appendix E has their detailed definitions.

Commands' semantics is also similar. $(c, \Omega_1) \hookrightarrow_C (r, \Omega_2)$ means running command c from stack Ω_1 results in a return status r which is either `cont` or `ret e` and a new stack Ω_2 . Note it no longer generates a history as the purpose of the deep simulation is merely computing K_0 and K_1 values. Also, it executes **normal**, **atomic**, **plocal1** and **box1** functions while ignores **plocal2** and **box2**. The detailed rules are in appendix E.

6.2 Correctness of Deep Simulation

Similar to shallow simulation, we define a refinement function \mathcal{Q} that converts R-values to symbolic R-values. This time we only turn all values with `pub2`, `pvt2` and `plc2` security levels into **SVsym**. This function is also overridden to apply on stacks.

$$\begin{array}{ccc} \mathcal{Q}(\Omega_1) & \xrightarrow[r']{c} & \mathcal{Q}(\Omega_2) \\ \vdots & & \vdots \\ \Omega_1 & \xrightarrow[\beta, r]{c} & \Omega_2 \end{array}$$

Figure 8: Deep Semantics Refinement

Theorem 5 *If $(c, \Omega_1) \rightarrow_C (r, \Omega_2, \beta)$ then there exists r' where*

$$(c, \mathcal{Q}(\Omega_1)) \hookrightarrow_C (r', \mathcal{Q}(\Omega_2)) \text{ and } \llbracket r' \rrbracket_R^{\Omega_2} = r.$$

The proof is very similar to theorem 1 as their semantics are similar except that deep semantics is a bit simpler.

6.3 Correctness of Distribution

Sec. 5.3 shows the distributed program, and Sec. 6.1 evaluates all dependent data that are needed in the `sync` steps. We now prove these data are sufficient for the distributed program to terminate with the same behavior as if running without distribution.

Given a sequential program $c_1; \dots; c_n$ generated by the shallow simulation and an initial stack Ω_0 , suppose the program runs like $(c_i, \Omega_{i-1}) \rightarrow_C (\text{cont}, \Omega_i, \beta_i)$ without cutting. We focus on the t -th chunk and align its execution like in Figure 9. Each command c_j corresponds to two stacks Ω_j^{pre} and Ω_j^{post} . If $\text{chunk}(j) = t$, we first run `syncj` to add external dependent values into $\Omega_{j-1}^{\text{post}}$, then run c_j ; Otherwise, the stack remains unchanged.

$$\begin{array}{ccccc} \Omega_{j-1}^{\text{post}} & \xrightarrow{\text{sync}_j} & \Omega_j^{\text{pre}} & \xrightarrow[\text{cont}, \beta_j]{c_j} & \Omega_j^{\text{post}} & \text{---} & \Omega_{j+1}^{\text{pre}} & \text{---} & \Omega_{j+1}^{\text{post}} \\ \Omega_{j-1} & \xrightarrow[\text{cont}, \beta_j]{c_j} & \Omega_j & \xrightarrow{c_{j+1}} & \Omega_{j+1} \end{array}$$

Figure 9: Execution in a Chunk

We need to prove that the following theorem holds, then theorem 4 ensures that c_j terminates with the same behavior.

Theorem 6 $\forall j \in \text{chunk}^{-1}(t), \Omega_j^{\text{pre}} \approx_S \Omega_{j-1}$ where $S = \text{DEP}(-, j)$.

In appendix F, we prove this theorem by first defining a set of L-values ϕ_j^t that are internally dependent in the t -th chunk after c_j , then prove the stack in the uncut execution is always similar to the t -th chunk’s local stack up to this set. So the internal dependency is handled, we then use theorem 5 to handle the external dependency happening in `sync`, combining them together proves the theorem.

7 Evaluation

7.1 Evaluation Setup and Metrics

We use three metrics to evaluate the effectiveness of our tool:

- **Compilation Time** is the time used for shallow simulation, data dependency analysis, and efficient-cut search. PBO solving during the compilation is done via the Gurobi [GO20] PBO solver. We do not include the time cost of deep simulation in the compilation time, but deep simulation takes less than 0.5% of the compilation time for all examples.
- **Effective Ratio** of a distributed zero knowledge program is defined as the quotient of the cost of the sequential execution (total costs of all instructions) by the cost of the distributed execution (as the objective function in Sec. 5). When a program is distributed to κ chunks, the closer the effective ratio is to κ , the better the distribution scheme is. We note that the optimal effective ratio may be smaller than κ , depending on the statement.
- **Execution Time** refers to the end-to-end performance. We compile Ou code to a VOLE-based ZK implementation (EMP [WMK16]) and measure the execution time of the resulting distributed programs. We opt to use VOLE-based ZK backend as one of our motivations is to work with large-scale applications. This choice allows us to handle large statements without requiring costly machines with large memory. In addition, our implementation can be easily extended to support a rich set of ZK protocols, since EMP can be connected to SIEVE-IR [Pro22].

To evaluate our tool, we focus on two programs with different features. All implementations in our language Ou can be found in the supplementary material. As the effectiveness of Lian also depends on the structure of the input program, we further provide visualization for the program structure of gradient descent and Merkle tree in Appendix G.

1. **Gradient descent (GD)** is an optimization algorithm for finding the optimal model in machine learning. It works by iteratively updating the parameters of a model in the direction of the negative gradient of the cost function concerning those parameters. We implement gradient decent for logistic regression with 10 features.
2. **Merkle tree (MT)** is a type of data structure used to verify the integrity of large amounts of data. It breaks the data into smaller blocks, then recursively hashes each pair of blocks to create a new set of parent nodes using a cryptographic hash function. By proving the correct computation of the Merkle tree, the prover can prove they have the original data without revealing the actual data to the verifier [Mer87].

7.2 Evaluation of the Partition Quality

First, we demonstrate the changes in the effective ratio of the distribution scheme found by Lian for GD and MT program changes with the size of programs and the number of cuts. GD takes a set $\{(\mathbf{x}, y)\}$ as input, where \mathbf{x} is a vector of size M . Throughout this section, we fix $|\{(\mathbf{x}, y)\}|$ to be 20 and $|\mathbf{x}|$ to be 10.

The effective ratio of Lian is affected by the length of the program being distributed. The length of GD and MT programs are decided by the number of iterations and the number of input blocks, respectively. As shown in Figure 10b and Figure 11b, Lian can find the distribution scheme that achieves the almost optimal effective ratio for large programs for both GD and MT cases.

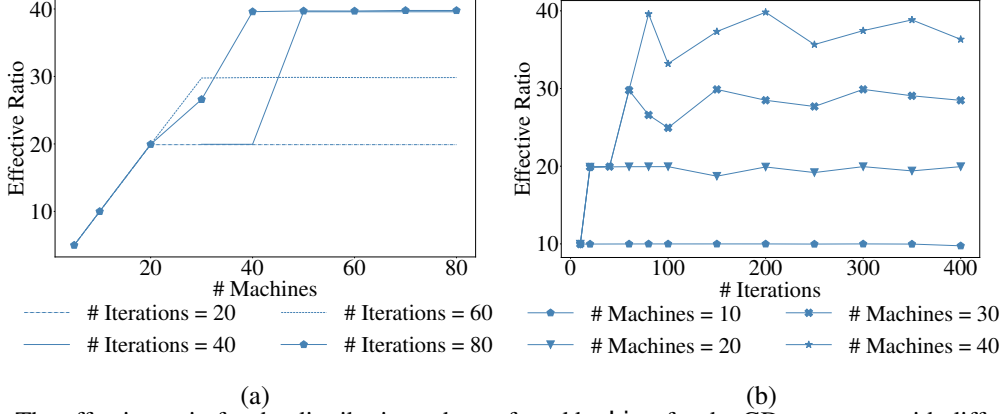


Figure 10: The effective ratio for the distribution scheme found by Lian for the GD programs with different numbers of iterations and various numbers of machines (κ). With the increase in κ , the effective ratio will reach the limit. A small-size program with a large κ will make the effective ratio of the distribution scheme much smaller than κ .

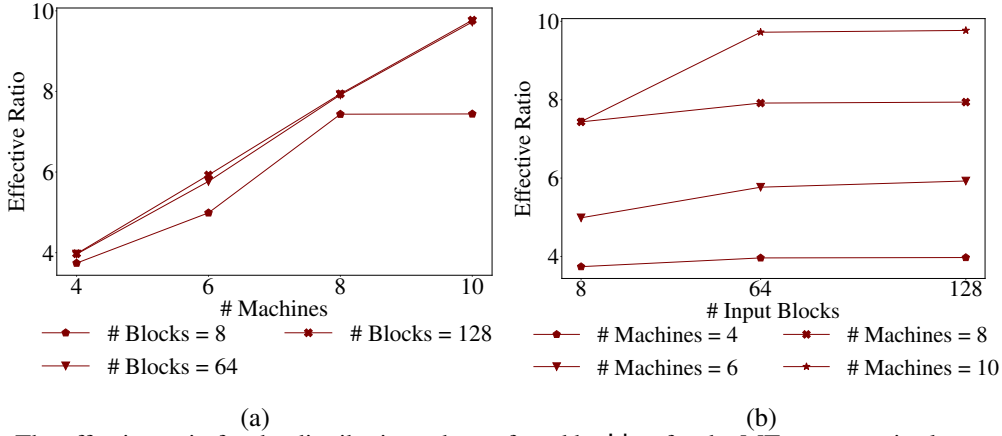


Figure 11: The effective ratio for the distribution scheme found by Lian for the MT programs is close to the numbers of machines (κ) with different numbers of input blocks and κ .

Figure 10a and Figure 11a demonstrate that the effective ratio will increase with the number of machines rising when the number of partitions is relatively small. Moreover, the results present an upper bound of the effective ratio of the distribution scheme for a given program. For example, the effective ratio for the GD programs involving 80 iterations reaches 40 and stays unchanged regardless of the increasing number of machines. Nevertheless, the results show that our tool can utilize a high degree of parallelization.

7.3 Effectiveness of Atomic Annotation

In this section, we study the impact of annotation to the compilation time and effective ratio. For all examples, we set the time limit of the PBO solver to 10 minutes. In the program for GD, there is a function called `update_all()`. This function updates the model's parameter using all data points and is called for each iteration. There is a function called `sha256_node()` that computes the hash of internal blocks in the MT program. We compare the compilation time and effective ratio of GD programs (or MT programs) with and without `update_all()` (or `sha256_node()`) being atomic while changing κ and the number of iterations.

The results of our evaluation show that the effects of adding atomic annotations depend on the structure of the programs. Figure 12 and Figure 13 show a significant rise in the compilation efficiency when adding annotation for both GD and MT programs. The compilation time decreases by 1X to 60X with κ and the number of iterations varying. The price of efficient compilation is a reduction in the effective ratio. The effective ratio for the annotated program is smaller than the unannotated one for GD programs: adding

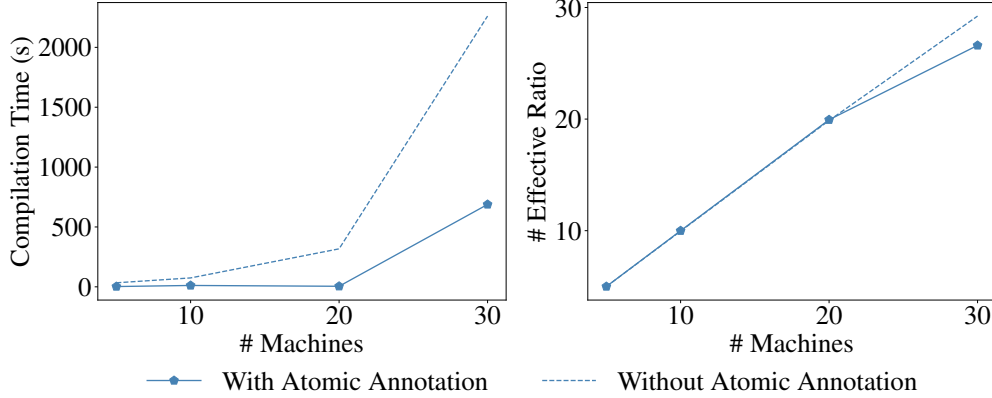


Figure 12: Compiling takes significantly greater time for GD programs when the function `update_all()` is not annotated with the atomic keyword. On the other hand, atomic annotation rules out some feasible solutions, making the distribution scheme slightly suboptimal.

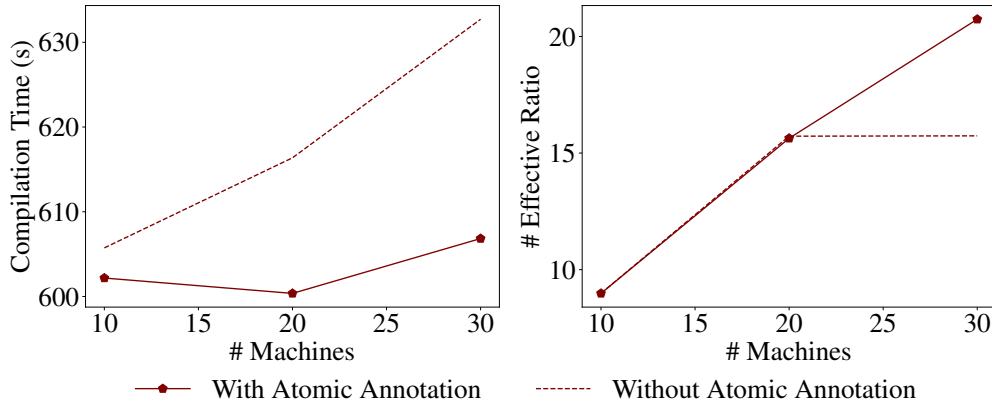


Figure 13: Adding annotation improves both the effectiveness of both compilation and the effective ratio of the MT program for 32-block input.

atomic annotation rules out a part of feasible solutions which may contain the optimum.

Sometimes, adding annotation could even improve the effective ratio, especially when the constrained time is far less than what is needed to search for an efficient partition. Figure 13 demonstrates an example of this case. We measure the effective ratio while setting the time limit of the PBO solver to 600 seconds. The effective ratio of the MT program for 32-block input with annotation is even better than the one without annotation when the number of cuts increases. One reason behind this result is the MT program structure. The solution found for the annotated program is also optimum for the unannotated one. Furthermore, with a part of the functions being atomic, the search space of the PBO solving essentially shrinks. For readers who are interested in additional information, we have included the results of MT programs for various input sizes in the appendix.

7.4 End-to-End Running Time

Our implementation is end-to-end: it can compile an Ou program to κ pieces of EMP circuits. Most of the compiler’s implementation is not tied to EMP, so it is flexible to connect with other backends.

We test the running time of the distributed circuits generated from GD and MT. The testing machine is an AWS m5.large instance with 2vCPU and 8GB memory. For GD, we fix the number of dataset to be 10, the size of each dataset to be 100 and iterate 200 rounds. For MT, we fix the data length to be 256. Table 1 shows with different number of machines, the slowest part’s running time and the overall speedup comparing to uncut execution time. We also list the effective ratio estimated by the solver. For GD, the effective ratio

κ	1	5	10	20	40
runtime GD	681.55 s	147.91 s	71.23 s	38.06 s	21.90 s
runtime speedup GD	1	4.6	9.57	17.9	31.12
effective ratio GD	1	4.99	9.95	19.9	39.8
runtime MT	40.23 s	11.08 s	7.63 s	6.5 s	5.25 s
runtime speedup MT	1	3.63	5.27	6.19	7.66
effective ratio MT	1	4.95	9.75	19.47	38.63

Table 1: End-to-end running time of our framework.

approximates the real speedup very well; For MT, the effective ratio becomes inaccurate as each chunk’s runtime reduces. We think the reason is mainly because the underlying circuits need a constant setup time (around 4 seconds). After subtracting it, the speedup becomes close to the effective ratio.

Acknowledgement

This work is supported by NSF awards CCF-2106845, CCF-2131476, CCF-1763399, CNS-2016240, CNS-2236819, CCF-2219995, DARPA under Contract No. HR001120C0087, and the Office of Naval Research (ONR) of the United States Department of Defense through an National Defense Science and Engineering Graduate (NDSEG) Fellowship. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- [ACBM08] Elli Androulaki, Seung Geol Choi, Steven M. Bellovin, and Tal Malkin. Reputation systems for anonymous networks. In *PETS 2008*, volume 5134 of *LNCS*, pages 202–218. Springer, Heidelberg, July 2008.
- [AGM18] Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 643–673. Springer, Heidelberg, August 2018.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- [BCC⁺16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Heidelberg, May 2016.
- [BCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013.
- [BCG⁺14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014.

- [BCG⁺20] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. ZEXE: Enabling decentralized private computation. In *2020 IEEE Symposium on Security and Privacy*, pages 947–964. IEEE Computer Society Press, May 2020.
- [BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *USENIX Security 2014*, pages 781–796. USENIX Association, August 2014.
- [BLW08] Dan Bogdanov, Sven Laur, and Jan Willemsen. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS 2008*, volume 5283 of *LNCS*, pages 192–206. Springer, Heidelberg, October 2008.
- [BMRS20] Joseph Bonneau, Izaak Meckler, Vanishree Rao, and Evan Shapiro. Coda: Decentralized cryptocurrency at scale. Cryptology ePrint Archive, Report 2020/352, 2020. <https://eprint.iacr.org/2020/352>.
- [BMRS21] Carsten Baum, Alex J. Malozemoff, Marc B. Rosen, and Peter Scholl. Mac’n’cheese: Zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 92–122, Virtual Event, August 2021. Springer, Heidelberg.
- [BMS⁺16] Aydın Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. Recent Advances in Graph Partitioning. *Algorithm Engineering*, pages 117–158, 2016.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [CFF⁺21] Matteo Campanelli, Antonio Faonio, Dario Fiore, Anaïs Querol, and Hadrián Rodríguez. Lunar: A toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 3–33. Springer, Heidelberg, December 2021.
- [CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019.
- [Cla13] Arthur C. Clarke. *Profiles of the Future*. Hachette UK, 2013. Originally 1962.
- [CR93] Sunil Chopra and Mendu R Rao. The partition problem. *Mathematical programming*, 59(1):87–115, 1993.
- [DILO22] Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Improving line-point zero knowledge: Two multiplications for the price of one. Cryptology ePrint Archive, Report 2022/552, 2022. <https://eprint.iacr.org/2022/552>.
- [DIO20] Samuel Dittmer, Yuval Ishai, and Rafail Ostrovsky. Line-point zero knowledge and its applications. Cryptology ePrint Archive, Report 2020/1446, 2020. <https://eprint.iacr.org/2020/1446>.
- [FPS⁺18] Jonathan Frankle, Sunoo Park, Daniel Shaar, Shafi Goldwasser, and Daniel J. Weitzner. Practical accountability of secret processes. In *USENIX Security 2018*, pages 657–674. USENIX Association, August 2018.

- [GM82] Joseph A. Goguen and José Meseguer. Security policies and security models. In *1982 IEEE Symposium on Security and Privacy, Oakland, CA, USA, April 26-28, 1982*, pages 11–20. IEEE Computer Society, 1982.
- [GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, October 1986.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software Protection and Simulation on Oblivious RAMs. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.
- [GO20] L Gurobi Optimization. Gurobi optimizer reference manual (2020), 2020.
- [Gol09] Oded Goldreich. *Foundations of cryptography: volume 2, basic applications*. Cambridge university press, 2009.
- [Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, December 2010.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- [HK20] David Heath and Vladimir Kolesnikov. Stacked garbling for disjunctive zero-knowledge proofs. In *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 569–598. Springer, Heidelberg, May 2020.
- [JKO13] Marek Jawurek, Florian Kerschbaum, and Claudio Orlandi. Zero-knowledge using garbled circuits: how to prove non-algebraic statements efficiently. In *ACM CCS 2013*, pages 955–966. ACM Press, November 2013.
- [KsMB13] Benjamin Kreuter, abhi shelat, Benjamin Mood, and Kevin R. B. Butler. PCF: A portable circuit format for scalable two-party secure computation. In *USENIX Security 2013*, pages 321–336. USENIX Association, August 2013.
- [LWN⁺15] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. OblivM: A programming framework for secure computation. In *2015 IEEE Symposium on Security and Privacy*, pages 359–376. IEEE Computer Society Press, May 2015.
- [Mer87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In *CRYPTO*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378. Springer, 1987.
- [PF79] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *Journal of the ACM (JACM)*, 26(2):361–381, 1979.
- [Pro22] The DARPA SIEVE Program. Sieve intermediate representation, 2022.
- [Set20] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 704–737. Springer, Heidelberg, August 2020.
- [SM03] Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, 2003.

- [WJB⁺17] Riad S Wahby, Ye Ji, Andrew J Blumberg, Abhi Shelat, Justin Thaler, Michael Walfish, and Thomas Wies. Full accounting for verifiable outsourcing. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 2071–2086, 2017.
- [WMK16] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient MultiParty Computation Toolkit. <https://github.com/emp-toolkit>, 2016.
- [WTs⁺18] Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018.
- [WYKW21] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *2021 IEEE Symposium on Security and Privacy*, pages 1074–1091. IEEE Computer Society Press, May 2021.
- [WZC⁺18] Howard Wu, Wenting Zheng, Alessandro Chiesa, Raluca Ada Popa, and Ion Stoica. DIZK: A distributed zero knowledge proof system. In *USENIX Security 2018*, pages 675–692. USENIX Association, August 2018.
- [XZC⁺22] Tiancheng Xie, Jiaheng Zhang, Zerui Cheng, Fan Zhang, Yupeng Zhang, Yongzheng Jia, Dan Boneh, and Dawn Song. zkBridge: Trustless cross-chain bridges made practical. In *ACM CCS 2022*, pages 3003–3017. ACM Press, November 2022.
- [YHKD22] Yibin Yang, David Heath, Vladimir Kolesnikov, and David Devecsery. EZEE: Epoch parallel zero knowledge for ANSI C. Cryptology ePrint Archive, Report 2022/811, 2022. <https://eprint.iacr.org/2022/811>.
- [zkp] ZKProof Community Reference. <https://docs.zkproof.org/reference.pdf>.
- [ZXZS20] Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. Transparent polynomial delegation and its applications to zero knowledge proof. In *2020 IEEE Symposium on Security and Privacy*, pages 859–876. IEEE Computer Society Press, May 2020.

A Example

Figure 14 shows a code snippet of the Merkle tree we used in the benchmark. The prover uses a binary Merkle tree to compute the hash of N private data blocks. The `merkle` function computes the hash of a segment of data blocks. Then the verifier calls `verify` to verify the hash. When the tree is a leaf node, `merkle` calls `sha256_leaf` to fetch a data block and hash it; When the tree is a non-leaf node, `merkle` recursively computes its two subtrees’ hashes, then use `sha256_node` to compute the merged hash. Suppose the user wants to cut it into 2 fragments, and there are 4 data blocks.

Here we annotate `sha256_node` and `sha256_leaf` as **atomic**, and `load_block` as **plocal1**. So during the shallow simulation, `sha256_node` and `sha256_leaf` will not be unfolded, and `load_block` will be ignored during the shallow simulation and executed during the deep simulation.

Shallow simulation executes the program with K_0 knowledge, yielding a sequence as Figure 15. The recursive function `merkle` is unfolded. Each variable is annotated with a unique number to distinguish it from others with the same name. Line 1-5 computes the hash of data block 0 and 1, line 6-11 computes the hash of data block 2 and 3, line 13-14 merges them, then line 15 verifies it.

Notice in Figure 14 `merkle`, the recursion condition is a **pub0** expression. This is necessary for the unfolding to proceed during the shallow simulation phase. If there is any non-**pub0** expression involved in deciding the control flow, then 1) the shallow simulation does not have enough knowledge to evaluate

```

1 #define N 4
2 /* compute merkle hash of block[left, ..., right-1] */
3 pvt1 int[8] merkle(int left, int right) {
4     if (left + 1 >= right) {
5         pvt1 int[8] hash = sha256_leaf(left);
6         return hash;
7     }
8     else {
9         int mid = (left + right) / 2;
10        pvt1 int[8] hashL = merkle(left, mid);
11        pvt1 int[8] hashR = merkle(mid, right);
12        pvt1 int[8] hash = sha256_node(&hashL, &hashR);
13        return hash;
14    }
15 }
16
17 plocal1 plc1 int[64] load_block(int block_id) { ... }
18
19 /* compute the hash of a block */
20 pvt1 int[8] sha256_block(pvt1 int[64] &data) { ... }
21
22 atomic void verify(pvt1 int[8] *hash) { ... }
23
24 /* compute the hash of a data block */
25 atomic pvt1 int[8] sha256_leaf(int block_id) {
26     /* prover loads and commits a block */
27     plc1 int[64] plc_block = load_block(block_id);
28     pvt1 int[64] block = {0};
29     for int i = 0; i < 64; i = i + 1; {
30         block[i] = commit(plc_block[i]);
31     }
32     pvt1 int[8] hash = sha256_block(block);
33     return hash;
34 }
35
36 /* compute the hash of two merged hashes */
37 atomic pvt1 int[8] sha256_node(
38     pvt1 int[8] *hashL, pvt1 int[8] *hashR) {
39     pvt1 int[64] block = {0};
40     /* merge two hashes */
41     for int i = 0; i < 8; i = i + 1; {
42         block[i] = hashL->[i];
43         block[i+8] = hashR->[i];
44     }
45     for int i = 16; i < 64; i = i + 1; {
46         block[i] = 0;
47     }
48     /* compute the hash of the merged results */
49     pvt1 int[8] hash = sha256_block(block);
50     return hash;
51 }
52
53 unit main() {
54     pvt1 int[8] hash = merkle(0, N);
55     verify(&hash);
56     return;
57 }

```

Figure 14: Merkle Tree Example

```

1 pvt1 int[8] hash1 = sha256_leaf(0);
2 pvt1 int[8] hashL1 = hash1;
3 pvt1 int[8] hash3 = sha256_leaf(1);
4 pvt1 int[8] hashR0 = hash3;
5 pvt1 int[8] hash5 = sha256_node(&hashL1, &hashR0);
6 pvt1 int[8] hashL0 = hash5;
7 pvt1 int[8] hash7 = sha256_leaf(2);
8 pvt1 int[8] hashL3 = hash7;
9 pvt1 int[8] hash9 = sha256_leaf(3);
10 pvt1 int[8] hashR3 = hash9;
11 pvt1 int[8] hash11 = sha256_node(&hashL3, &hashR3);
12 pvt1 int[8] hashR2 = hash11;
13 pvt1 int[8] hash13 = sha256_node(&hashL0, &hashR2);
14 pvt1 int[8] hash0 = hash13;
15 verify(&hash0);

```

Figure 15: Merkle Tree after Shallow Simulation

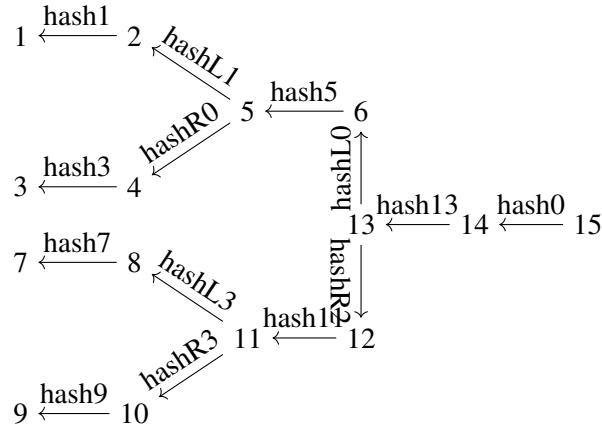


Figure 16: Merkle Tree Dependency Graph

that expression so it can not decide which branch to take, and 2) even if it uses the prover’s knowledge to evaluate that expression, the information will be leaked to the verifier via the generated programs.

Then live variable analysis generates a dependency graph as in Figure 16. The ILP/PBO solver decides to cut it between lines 5 and 6, so the only dependency is `hash5`.

Now the compiler can already generate two programs as in Figure 17. We only show the produce and sync functions here. Note `produce_5` in chunk 1 records a digest of all variables that are produced by Figure 15’s line 5 for consistency check in the end. Here it is `hash5`. Similarly, `sync_6` records a digest of all variables dependent by Figure 15’s line 6 after assigning values to them. But we do not know `hash5`’s value yet, so we mark it using a question mark. The code can already be distributed to individual verifiers, who will count on the provers to supply encrypted `hash5` during the runtime. But the provers need to know these values before interacting with the verifiers. So they perform deep simulation on Figure 15’s program based on the secret input and find out `hash5`’s value when reaching line 5.

Here the type system again guarantees the deep simulation can proceed. For example, suppose there is any K_2 value that is dependent between two chunks, and it relies on the verifier’s randomness, then the prover at the compile time can not guess this value, thus it can not proceed.

In the end, chunk 1 and chunk 2 can run in parallel without any communication. When both chunks’ execution finish, they will run a consistency check together to make sure the `hash5` used in chunk 2 is indeed produced by chunk 1’s execution.

Here we can see the live variable analysis and the cut searching should find the minimal dependency

```

1 /* chunk 1 */
2 pvt1 int[8] hash1 = sha256_leaf(0);
3 pvt1 int[8] hashL1 = hash1;
4 pvt1 int[8] hash3 = sha256_leaf(1);
5 pvt1 int[8] hashR0 = hash3;
6 pvt1 int[8] hash5 = sha256_node(&hashL1, &hashR0);
7 produce_5();
8 consistency_check();
9
10 void produce_5() {
11     add_out_digest(2, hash5);
12 }

```

```

1 /* chunk 2 */
2 pvt1 int[8] hash5;
3 sync_6();
4 pvt1 int[8] hashL0 = hash5;
5 pvt1 int[8] hash7 = sha256_leaf(2);
6 pvt1 int[8] hashL3 = hash7;
7 pvt1 int[8] hash9 = sha256_leaf(3);
8 pvt1 int[8] hashR3 = hash9;
9 pvt1 int[8] hash11 = sha256_node(&hashL3, &hashR3);
10 pvt1 int[8] hashR2 = hash11;
11 pvt1 int[8] hash13 = sha256_node(&hashL0, &hashR2);
12 pvt1 int[8] hash0 = hash13;
13 verify(&hash0);
14 consistency_check();
15
16 void sync_6() {
17     hash5 = ?;
18     add_in_digest(1, hash5);
19 }

```

Figure 17: Merkle Tree Generated Programs

between the two chunks that guarantees the safe execution of both. Suppose the live variable analysis fails to determine that `hash5` is the dependency between chunk 1 and chunk 2, then chunk 2's line 4 will have undefined behavior. Suppose the cut searching algorithm considers `hash1`, `hash3` and `hash5` are all needed by chunk2, then we need to commit more data in `sync_6`, which will slow down the runtime performance.

$$\begin{array}{ll}
\llbracket x \rrbracket_L^\Omega = \mathbf{Vvar} \text{ find}(\Omega, x) & \llbracket \ell n \rrbracket_R^\Omega = \mathbf{Vint} \ell n \\
\llbracket \epsilon[e] \rrbracket_L^\Omega = \mathbf{Vindex} \llbracket \epsilon \rrbracket_L^\Omega \ell n & \llbracket \epsilon \rrbracket_R^\Omega = \text{load}(\Omega, \llbracket \epsilon \rrbracket_L^\Omega) \\
\text{where } \llbracket e \rrbracket_R^\Omega = \mathbf{Vint} \ell n & \llbracket \& \epsilon \rrbracket_R^\Omega = \mathbf{Vref} \llbracket \epsilon \rrbracket_L^\Omega \\
\llbracket \epsilon.l \rrbracket_L^\Omega = \mathbf{Vfield} \llbracket \epsilon \rrbracket_L^\Omega l & \llbracket \ell \{e_0, \dots, e_{n-1}\} \rrbracket_R^\Omega = \mathbf{Varray} \ell \cup \{i \mapsto \llbracket e_i \rrbracket_R^\Omega\} \\
\llbracket * e \rrbracket_L^\Omega = \mu \text{ where } \llbracket e \rrbracket_R^\Omega = \mathbf{Vref} \mu & \llbracket f(e_1, \dots, e_n) \rrbracket_R^\Omega = (f)(\llbracket e_1 \rrbracket_R^\Omega, \dots, \llbracket e_n \rrbracket_R^\Omega)
\end{array}$$

$$\boxed{(c, \Omega) \rightarrow_C (r, \Omega', \beta)}$$

$$\frac{}{(\epsilon = e, \Omega) \rightarrow_C (\text{cont}, \text{store}(\Omega, \llbracket \epsilon \rrbracket_L^\Omega, \llbracket e \rrbracket_R^\Omega), \langle \rangle) \text{ (assert } e, \Omega) \rightarrow_C (\text{cont}, \Omega, \llbracket e \rrbracket_R^\Omega)}$$

$$\frac{}{(\text{return } e, \Omega) \rightarrow_C (\text{ret } \llbracket e \rrbracket_R^\Omega, \Omega, \langle \rangle)} \quad \frac{\llbracket e \rrbracket_R^\Omega = \mathbf{Vbool} \ell \text{ true} \quad (c_1, \Omega) \rightarrow_C (r, \Omega', \beta)}{(\text{if } e \text{ then } c_1 \text{ else } c_2, \Omega) \rightarrow_C (r, \Omega', \beta)}$$

$$\frac{\llbracket e \rrbracket_R^\Omega = \mathbf{Vbool} \ell \text{ true}}{(c; \text{while } e \text{ do } c, \Omega) \rightarrow_C (r, \Omega', \beta)} \quad \frac{\llbracket e \rrbracket_R^\Omega = \mathbf{Vbool} \ell \text{ false}}{(\text{while } e \text{ do } c, \Omega) \rightarrow_C (\text{cont}, \Omega, \langle \rangle)}$$

$$\begin{array}{c}
\text{new}(x) = \alpha \\
f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \\
\forall i \in [1, n], \llbracket e \rrbracket_R^\Omega = v_i \wedge \text{new}(x_i) = \alpha_i \\
(c, \bigcup \{\alpha_i \mapsto v_i\} \cdot \Omega) \rightarrow_C (\text{ret } v, \omega \cdot \Omega', \beta) \\
\hline
(\tau x = f(e_1, \dots, e_n), \Omega) \rightarrow_C (\text{cont}, \Omega'[\alpha \mapsto v], \beta) \\
\hline
\text{new}(x) = \alpha \quad (c_1, \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta) \\
\hline
(\tau x = e, \Omega) \rightarrow_C (\text{cont}, \Omega[\alpha \mapsto \llbracket e \rrbracket_R^\Omega], \langle \rangle) \quad (c_1; c_2, \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta) \\
\hline
(c_1, \Omega) \rightarrow_C (\text{cont}, \Omega_1, \beta_1) \quad (c_2, \Omega_1) \rightarrow_C (r, \Omega_2, \beta_2) \\
\hline
(c_1; c_2, \Omega) \rightarrow_C (r, \Omega_2, \beta_1 \cdot \beta_2)
\end{array}$$

Figure 18: Dynamic Semantics

B Detailed Semantics

The dynamic semantics uses the following values:

$$\begin{array}{l}
v ::= \mathbf{Vunit} \mid \mathbf{Vint} \ell \text{ num} \\
\quad \mid \mathbf{Vfloat} \ell \text{ num} \mid \mathbf{Vbool} \ell \text{ bool} \\
\quad \mid \mathbf{Vnull} \mid \mathbf{Vref} \mu \\
\quad \mid \mathbf{Varray} \ell \text{ array} \mid \mathbf{Vstruct} \text{ map} \\
\mu ::= \mathbf{Vvar} \alpha \mid \mathbf{Vindex} \mu \ell \text{ num} \mid \mathbf{Vfield} \mu \text{ label}
\end{array}$$

Figure 18 defines its full semantics.

Fact 2

$$((c_1; c_2); c_3, \Omega) \rightarrow_C (r, \Omega', \beta) \iff (c_1; (c_2; c_3), \Omega) \rightarrow_C (r, \Omega', \beta)$$

Proof.

1. Suppose that $(c_1, \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta)$. Then $(c_1; c_2, \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta)$, so $((c_1; c_2); c_3, \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta)$. It follows as well that $(c_1; (c_2; c_3), \Omega) \rightarrow_C (\text{ret } v, \Omega_1, \beta)$.

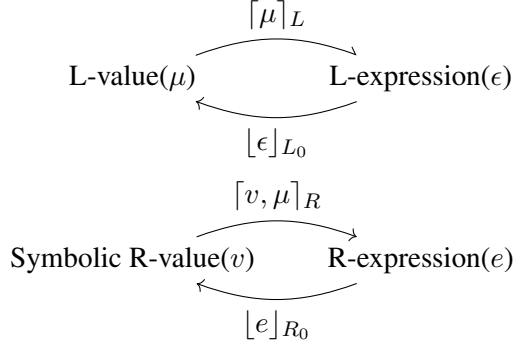


Figure 19: Converting symbolic values and expressions.

2. Suppose both that $(c_1, \Omega) \rightarrow_C (\text{cont}, \Omega_1, \beta_1)$ and $(c_2, \Omega_1) \rightarrow_C (\text{ret } v, \Omega_2, \beta_2)$. Then $(c_1; c_2, \Omega) \rightarrow_C (\text{ret } v, \Omega_2, \beta_1 \cdot \beta_2)$, so $((c_1; c_2)c_3, \Omega) \rightarrow_C (\text{ret } v, \Omega_2, \beta_1 \cdot \beta_2)$. Also $(c_2; c_3, \Omega_1) \rightarrow_C (\text{ret } v, \Omega_2, \beta_2)$, so $(c_1; (c_2; c_3), \Omega) \rightarrow_C (\text{ret } v, \Omega_2, \beta_1 \cdot \beta_2)$.
3. Suppose $(c_1, \Omega) \rightarrow_C (\text{cont}, \Omega_1, \beta_1)$, $(c_2, \Omega_1) \rightarrow_C (\text{cont}, \Omega_2, \beta_2)$ and $(c_3, \Omega_2) \rightarrow_C (r, \Omega_3, \beta_3)$. Then $(c_1; c_2, \Omega) \rightarrow_C (\text{cont}, \Omega_2, \beta_1 \cdot \beta_2)$, so $((c_1; c_2); c_3, \Omega) \rightarrow_C (r, \Omega_3, \beta_1 \cdot \beta_2 \cdot \beta_3)$. It follows as well that $(c_2; c_3, \Omega_1) \rightarrow_C (r, \Omega_3, \beta_2 \cdot \beta_3)$, so $(c_1; (c_2; c_3), \Omega) \rightarrow_C (r, \Omega_3, \beta_1 \cdot \beta_2 \cdot \beta_3)$.

C Detailed Shallow Simulation

C.1 Shallow Semantics

Here are some rules not listed in the main content.

$$\frac{(c_1, \Omega) \rightsquigarrow_C (\text{ret } e, \Omega_1, h)}{(c_1; c_2, \Omega) \rightsquigarrow_C (\text{ret } e, \Omega_1, h)}$$

$$\frac{\llbracket e \rrbracket_{R_0}^\Omega = \text{pub0 false} \quad (c_2, \Omega) \rightsquigarrow_C (r, \Omega', h)}{(\mathbf{if } e \mathbf{ then } c_1 \mathbf{ else } c_2, \Omega) \rightsquigarrow_C (r, \Omega', h)}$$

$$\frac{\llbracket e \rrbracket_{R_0}^\Omega = \text{pub0 true} \quad (c; \mathbf{while } e \mathbf{ do } c, \Omega) \rightsquigarrow_C (r, \Omega', h)}{(\mathbf{while } e \mathbf{ do } c, \Omega) \rightsquigarrow_C (r, \Omega', h)}$$

$$\frac{\llbracket e \rrbracket_{R_0}^\Omega = \text{pub0 false}}{(\mathbf{while } e \mathbf{ do } c, \Omega) \rightsquigarrow_C (\text{cont}, \Omega, \cdot)}$$

Here's the dynamic semantics of the two new commands:

$$\frac{}{(\mathbf{push}, \Omega) \rightarrow_C (\text{cont}, \emptyset \cdot \Omega, \langle \rangle)}$$

$$\frac{\llbracket e \rrbracket_R^{\omega \cdot \Omega} = v}{(\mathbf{pop } \tau \alpha = e, \omega \cdot \Omega) \rightarrow_C (\text{cont}, \Omega[\alpha \mapsto v], \langle \rangle)}$$

As with the dynamic semantics, we can show that command concatenation is associative.

Fact 3

$$S((c_1; c_2); c_3, \Omega) \rightsquigarrow_C (r, \Omega', h) \iff (c_1; (c_2; c_3), \Omega) \rightsquigarrow_C (r, \Omega', h).$$

The proof is similar to Fact 2.

We say an R-expression e (or L-expression ϵ) is *normalized* in a symbolic stack Ω , denoted as $\text{rnorm}(\Omega, e)$ (or $\text{lnorm}(\Omega, \epsilon)$), if all its sub-expressions cannot be further simplified.

$$\begin{array}{ll}
\text{lnorm}(\Omega, \alpha) & \\
\text{lnorm}(\Omega, \epsilon[e]) & \text{if } \text{lnorm}(\Omega, \epsilon) \wedge \text{rnorm}(\Omega, e) \\
\text{lnorm}(\Omega, \epsilon.l) & \text{if } \text{lnorm}(\Omega, \epsilon) \\
\\
\text{rnorm}(\Omega, \ell n) & \\
\text{rnorm}(\Omega, \& \epsilon) & \text{if } \text{lnorm}(\Omega, \epsilon) \\
\text{rnorm}(\Omega, \epsilon) & \text{if } \text{lnorm}(\Omega, \epsilon) \wedge \epsilon \notin \text{dom}[_]_{L_0} \\
& \text{or } \text{lnorm}(\Omega, \epsilon) \wedge \text{load}(\Omega, [\epsilon]_{L_0}) = \mathbf{SVsym} \\
\text{rnorm}(\Omega, \ell \{e_0, \dots\}) & \text{if } \forall i, \text{rnorm}(\Omega, e_i) \\
\text{rnorm}(\Omega, \{l_1 = e_1, \dots\}) & \text{if } \forall i, \text{rnorm}(\Omega, e_i) \\
\text{rnorm}(\Omega, f(e_1 \dots e_n)) & \text{if } f \text{ is a non-pub0 builtin function} \\
& \text{and } \forall i, \text{rnorm}(\Omega, e_i)
\end{array}$$

These definitions establish another helpful consequence.

Fact 4

$$\llbracket e \rrbracket_{R_0}^\Omega = e' \implies \text{rnorm}(\Omega, e') \text{ and } \llbracket \epsilon \rrbracket_{L_0}^\Omega = \epsilon' \implies \text{lnorm}(\Omega, \epsilon').$$

Fact 5 *No pointer dereference will appear in a generated program.*

C.2 Proof of Theorem 1

Before the proof itself, we first present some simple lemmas about interacting with the actual stack and the symbolic stack, and then show two further lemmas about the correctness of simplifying R-expressions and L-expressions.

Lemma 7 *If $[\epsilon]_{L_0} = \mu$, then for any stack Ω , $\llbracket \epsilon \rrbracket_L^\Omega = \mu$.*

The proof follows from structural induction on L-expressions.

Lemma 8 $\mathcal{R}(\text{load}(\Omega), \mu) = \text{load}(\mathcal{R}(\Omega), \mu)$.

Lemma 9 $\mathcal{R}(\text{store}(\Omega, \mu, v)) = \text{store}(\mathcal{R}(\Omega), \mu, \mathcal{R}(v))$.

Lemma 10 $\mathcal{R}(\Omega[\alpha \mapsto v]) = \mathcal{R}(\Omega)[\alpha \mapsto \mathcal{R}(v)]$.

Lemma 11 $\mathcal{R}(\omega \cdot \Omega) = \mathcal{R}(\omega) \cdot \mathcal{R}(\Omega)$.

Lemma 12 $\text{load}(\Omega, \mu) = v \implies \llbracket [\mathcal{R}(v), \mu]_R \rrbracket_R^\Omega = v$.

Lemma 13 *If $\text{rnorm}(\mathcal{R}(\Omega), e)$, then $[e]_{R_0} = \mathcal{R}(\llbracket e \rrbracket_R^\Omega)$.*

Lemma 14 *If $\text{lnorm}(\mathcal{R}(\Omega), \epsilon)$, then either $\epsilon \notin \text{dom}[_]_{L_0}$ or $[\epsilon]_{L_0} = \llbracket \epsilon \rrbracket_L^\Omega$.*

Lemma 15 $\llbracket e \rrbracket_R^\Omega = v \implies \exists e', \llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e' \wedge \llbracket e' \rrbracket_R^\Omega = v$.

Lemma 16 $\llbracket \epsilon \rrbracket_L^\Omega = \mu \implies \exists \epsilon', \llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon' \wedge \llbracket \epsilon' \rrbracket_L^\Omega = \mu$.

The two lemmas can be proven together by structural induction on R-expressions and L-expressions. We discuss a few tricky cases here.

R-expression ϵ . Suppose $\llbracket \epsilon \rrbracket_R^\Omega = v$, $\llbracket \epsilon \rrbracket_L^\Omega = \mu$ and $v = \text{load}(\Omega, \mu)$. Lemma 16 guarantees ϵ can be simplified to ϵ' where $\llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon'$ and $\llbracket \epsilon' \rrbracket_L^\Omega = \mu$. We want to prove there exists an expression e such that $\llbracket \epsilon \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e$ and $\llbracket e \rrbracket_R^\Omega = v$. Note it is equivalent to prove $\llbracket \text{sload}(\mathcal{R}(\Omega), \epsilon') \rrbracket_R^\Omega = v$ since $\llbracket \epsilon \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = \text{sload}(\mathcal{R}(\Omega), \llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)}) = \text{sload}(\mathcal{R}(\Omega), \epsilon')$. There are two cases:

1. $\epsilon' \notin \text{dom}[-]_{L_0}$. Then $\text{sload}(\mathcal{R}(\Omega), \epsilon') = \epsilon'$. We can prove the goal by $\llbracket \epsilon' \rrbracket_R^\Omega = \text{load}(\Omega, \llbracket \epsilon' \rrbracket_L^\Omega) = \text{load}(\Omega, \mu) = v$.
2. $\epsilon' \in \text{dom}[-]_{L_0}$. Lemma 7 tells us $\llbracket \epsilon' \rrbracket_{L_0}^\Omega = \llbracket \epsilon' \rrbracket_L^\Omega = \mu$. So $\text{sload}(\mathcal{R}(\Omega), \epsilon') = \lceil \text{load}(\mathcal{R}(\Omega), \mu), \mu \rceil_R$ which equals to $\lceil \mathcal{R}(v), \mu \rceil_R$ according to lemma 8. Then we finish the goal using lemma 12: $\llbracket \lceil \mathcal{R}(v), \mu \rceil_R \rrbracket_R^\Omega = v$.

R-expression $\& \epsilon$. Suppose $\llbracket \& \epsilon \rrbracket_R^\Omega = \mathbf{Vref} \mu$ and $\llbracket \epsilon \rrbracket_L^\Omega = \mu$. Lemma 16 shows there exists ϵ' such that $\llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon'$ and $\llbracket \epsilon' \rrbracket_L^\Omega = \mu$. We want to prove $\llbracket \& \epsilon \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = \& \epsilon'$ and $\llbracket \& \epsilon' \rrbracket_R^\Omega = \mathbf{Vref} \mu$. The first goal is direct from the definition. As for the second goal, by definition $\llbracket \& \epsilon' \rrbracket_R^\Omega = \mathbf{Vref} \llbracket \epsilon' \rrbracket_L^\Omega = \mathbf{Vref} \mu$.

L-expression $\epsilon[e]$. Suppose $\llbracket \epsilon[e] \rrbracket_L^\Omega = \mathbf{Vindex} \mu \ell n$, $\llbracket \epsilon \rrbracket_L^\Omega = \mu$ and $\llbracket e \rrbracket_R^\Omega = \mathbf{Vint} \ell n$. Lemma 15 and 16 guarantee e and ϵ can be simplified to e' and ϵ' , respectively, where $\llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'$, $\llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon'$, $\llbracket e' \rrbracket_R^\Omega = \mathbf{Vint} \ell n$ and $\llbracket \epsilon' \rrbracket_L^\Omega = \mu$. By definition we know $\llbracket \epsilon[e] \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon'[e']$. Also $\llbracket \epsilon'[e'] \rrbracket_L^\Omega = \mathbf{Vindex} \mu \ell n$.

L-expression $* e$. Suppose $\llbracket * e \rrbracket_L^\Omega = \mu$ and $\llbracket e \rrbracket_R^\Omega = \mathbf{Vref} \mu$. Lemma 16 shows there exists e' such that $\llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'$ and $\llbracket e' \rrbracket_R^\Omega = \mathbf{Vref} \mu$. Depending on whether e' can be simplified or not, there are two cases.

1. $e' = \& \epsilon$ for some ϵ . Then $\llbracket * e \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = \epsilon$ by definition. We also know $\llbracket \epsilon \rrbracket_L^\Omega = \mu$ because $\llbracket e' \rrbracket_R^\Omega = \mathbf{Vref} \mu$, so this case is proven.
2. e' can not be simplified. Then $\llbracket * e \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = * e'$ by definition. Also $\llbracket * e' \rrbracket_L^\Omega = \mu$ by definition.

We can finally prove theorem 1 by structural induction on commands.

assert e . Suppose $(\text{assert } e, \Omega) \rightarrow_C (\text{cont}, \Omega, v)$ and $v = \llbracket e \rrbracket_R^\Omega$. Lemma 15 guarantees there exists e' such that $\llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'$ and $\llbracket e' \rrbracket_R^\Omega = v$. We now have two subgoals: $(\text{assert } e, \mathcal{R}(\Omega)) \rightsquigarrow_C (\text{cont}, \mathcal{R}(\Omega), \text{assert } e')$ and $(\text{assert } e', \Omega) \rightarrow_C (\text{cont}, \Omega, v)$. Both can be proven using semantics definitions.

$x = f(e_1, \dots, e_n)$. Suppose $(x = f(e_1, \dots, e_n), \Omega) \rightarrow_C (\text{cont}, \Omega'[\alpha \mapsto v], \beta)$, $\text{new}(x) = \alpha$, $f \mapsto (\langle \tau_1 x_1, \dots, \tau_n x_n \rangle \rightarrow \tau, c) \in \Lambda^u$, $\forall i \in [1, n]$, $\llbracket e_i \rrbracket_R^\Omega = v_i \wedge \text{new}(x_i) = \alpha_i$ and $(c, \bigcup \{\alpha_i \mapsto v_i\} \cdot \Omega) \rightarrow_C (\text{ret } v, \omega \cdot \Omega', \beta)$. For simplicity, let $\Omega_c := \bigcup \{\alpha_i \mapsto v_i\} \cdot \Omega$. By lemma 15 we know there exist $e'_1 \dots e'_n$ such that $\forall i \in [1, n]$, $\llbracket e_i \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'_i \wedge \llbracket e'_i \rrbracket_R^\Omega = v_i$. Use lemma 13 we know $\llbracket e'_i \rrbracket_{R_0} = \mathcal{R}(v_i)$. Also lemma 10 shows $\bigcup \{\alpha_i \mapsto \mathcal{R}(v_i)\} \cdot \mathcal{R}(\Omega) = \mathcal{R}(\bigcup \{\alpha_i \mapsto v_i\} \cdot \Omega) = \mathcal{R}(\Omega_c)$. By induction hypothesis we know there exist h and e such that $(c, \mathcal{R}(\Omega_c)) \rightsquigarrow_C (\text{ret } e, \mathcal{R}(\omega \cdot \Omega'), h)$ and $(h, \Omega_c) \rightarrow_C (\text{cont}, \omega \cdot \Omega', \beta)$ and $\llbracket \text{ret } e \rrbracket_R^{\omega \cdot \Omega'} = v$. Let $h' := \mathbf{push}; \tau_1 \alpha_1 = e'_1; \dots; \tau_n \alpha_n = e'_n; h; \mathbf{pop} \tau \alpha = e$, by semantics' definition we can prove $(\tau x = f(e_1, \dots, e_n), \mathcal{R}(\Omega)) \rightsquigarrow_C (\text{cont}, \mathcal{R}(\Omega')[\alpha \mapsto v], h')$. Now we just need to prove $(h', \Omega) \rightarrow_C (\text{cont}, \Omega'[\alpha \mapsto v], \beta)$. Note for any $j \in [1, i]$, α_j will never appear in e_i , so evaluating e_i in $\emptyset \cdot \Omega$ is the same as evaluating it after defining x_1, \dots, x_{i-1} . So $(\tau_1 x_1 = e_1; \dots; \tau_n x_n = e_n, \emptyset \cdot \Omega) \rightarrow_C (\text{cont}, \Omega_c, \langle \rangle)$. The induction hypothesis already shows $(h, \Omega_c) \rightarrow_C (\text{cont}, \omega \cdot \Omega', \beta)$. By definition we

also know $(\mathbf{pop} \tau \alpha = e, \omega \cdot \Omega') \rightarrow_C (\text{cont}, \Omega'[\alpha \mapsto v], \langle \rangle)$. We can prove the goal by concatenating these commands together.

$\tau x = e$. Suppose $(\tau x = e, \Omega) \rightarrow_C (\text{cont}, \Omega[\alpha \mapsto v], \langle \rangle)$, $\text{new}(x) = \alpha$ and $v = \llbracket e \rrbracket_R^\Omega$. Lemma 15 guarantees there exists e' such that $\llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'$ and $\llbracket e' \rrbracket_R^\Omega = v$. By definition we know $(\tau x = e, \mathcal{R}(\Omega)) \rightsquigarrow_C (\text{cont}, \mathcal{R}(\Omega)[\alpha \mapsto \llbracket e' \rrbracket_{R_0}], \tau \alpha = e')$. So we only need to prove $\mathcal{R}(\Omega[\alpha \mapsto v]) = \mathcal{R}(\Omega)[\alpha \mapsto \llbracket e' \rrbracket_{R_0}]$ and $(\tau \alpha = e', \Omega) \rightarrow_C (\text{cont}, \Omega[\alpha \mapsto v], \langle \rangle)$. The second one is direct from the definition. As for the first one, by lemma 13 we know $\llbracket e' \rrbracket_{R_0} = \mathcal{R}(v)$, so $\mathcal{R}(\Omega)[\alpha \mapsto \llbracket e' \rrbracket_{R_0}] = \mathcal{R}(\Omega)[\alpha \mapsto \mathcal{R}(v)]$ which is equal to $\mathcal{R}(\Omega[\alpha \mapsto v])$ by lemma 10.

$\epsilon = e$. Suppose $(\epsilon = e, \Omega) \rightarrow_C (\text{cont}, \text{store}(\Omega, \mu, v), \langle \rangle)$, $\llbracket e \rrbracket_R^\Omega = v$ and $\llbracket \epsilon \rrbracket_L^\Omega = \mu$. Lemma 15 and 16 show there exist e' and e'' such that $\llbracket e \rrbracket_{R_0}^{\mathcal{R}(\Omega)} = e'$, $\llbracket \epsilon \rrbracket_{L_0}^{\mathcal{R}(\Omega)} = e''$, $\llbracket e' \rrbracket_R^\Omega = v$ and $\llbracket e'' \rrbracket_L^\Omega = \mu$. By definition we already know $(e' = e', \Omega) \rightarrow_C (\text{cont}, \text{store}(\Omega, \mu, v), \langle \rangle)$ and $(\epsilon = e, \mathcal{R}(\Omega)) \rightsquigarrow_C (\text{cont}, \text{sstore}(\mathcal{R}(\Omega), e', e''), e' = e')$, so we only need to prove $\text{sstore}(\mathcal{R}(\Omega), e', e'') = \mathcal{R}(\text{store}(\Omega, \mu, v))$. Depending on whether e'' can be fully evaluated or not, there are two cases.

1. If $e'' \notin \text{dom}[_]_{L_0}$, then $\text{sstore}(\mathcal{R}(\Omega), e', e'') = \mathcal{R}(\Omega)$ also μ has non-`pub0` information, so the type wellformedness guarantees the value pointed by it can not contain any `pub0`. So we know $\mathcal{R}(v)$ is full of `SVsym`, so $\mathcal{R}(\text{store}(\Omega, \mu, v)) = \mathcal{R}(\Omega)$.
2. If $e'' \in \text{dom}[_]_{L_0}$, then lemma 14 tells us $\llbracket e'' \rrbracket_{L_0} = \mu$. Also lemma 13 tells us $\llbracket e' \rrbracket_{R_0} = \mathcal{R}(v)$. So $\text{sstore}(\mathcal{R}(\Omega), e', e'') = \text{store}(\mathcal{R}(\Omega), \mu, \mathcal{R}(v))$ which equals to $\mathcal{R}(\text{store}(\Omega, \mu, v))$ by lemma 9.

Termination. Technically, we can not do structural induction on while loops as they may not terminate. But since in the hypothesis the command terminates with r and β , we can always find a ranking function on stacks that reduces as the program executes.

D Detailed Live Variable Analysis

Recall that $[_]_{L_0}$ is a partial function as it may fail when the L-expression refers to an array index with a non- K_0 access level. We can, however, define a total function $[_]_{L_0}^*$ to compute the smallest L-value with K_0 access level that includes the result of ϵ .

$$\begin{aligned} [\alpha]_{L_0}^* &= \mathbf{vvar} \alpha \\ [\epsilon[e]]_{L_0}^* &= \begin{cases} \mu & \text{if } [\epsilon[e]]_{L_0} = \mu \\ [\epsilon]_{L_0}^* & \text{otherwise} \end{cases} \\ [\epsilon.l]_{L_0}^* &= \begin{cases} \mu & \text{if } [\epsilon.l]_{L_0} = \mu \\ [\epsilon]_{L_0}^* & \text{otherwise} \end{cases} \end{aligned}$$

Due to Fact 5 we can ignore dereferencing pointers. Let the set $\text{prims}(\mu)$ contain all smallest L-values with K_0 access level that reside in μ . We use $\text{prims}(\mu)$ to define two functions $\text{Lread}(\epsilon)$ and $\text{Rread}(e)$ that collect all of the smallest K_0 L-values from ϵ and e 's subexpressions, respectively. We also define $\text{Pointed}(e)$ to collect all of the smallest L-values with K_0 access level that are pointed to by e 's subexpressions. It is used to analyze a sandboxed function's arguments, as every pointer passed into a sandboxed

function is potentially read. We omit its definition as it is near identical to $\text{Rread}(e)$ itself.

$$\begin{aligned}
\text{Lread}(\alpha) &= \emptyset \\
\text{Lread}(\epsilon[e]) &= \text{Lread}(\epsilon) \cup \text{Rread}(e) \\
\text{Lread}(\epsilon.l) &= \text{Lread}(\epsilon) \\
\\
\text{Rread}(\ell n) &= \emptyset \\
\text{Rread}(\& e) &= \emptyset \\
\text{Rread}(\epsilon) &= \text{prims}([\epsilon]_{L_0}^*) \cup \text{Lread}(\epsilon) \\
\text{Rread}(f(e_1, \dots)) &= \bigcup \text{Rread}(e_i) \\
\text{Rread}(\ell \{e_0, \dots\}) &= \bigcup \text{Rread}(e_i)
\end{aligned}$$

$$\begin{aligned}
\text{REF}(\tau \alpha = e) &= \text{Rread}(e) \\
\text{REF}(\tau \alpha = f(e_1, \dots)) &= \bigcup (\text{Rread}(e_i) \cup \text{Pointed}(e_i)) \\
\text{REF}(\tau \alpha = f(e_1, \dots) \mathbf{R} S_R \mathbf{W} S_W) &= \bigcup \text{Rread}(e_i) \cup S_R \\
\text{REF}(\mathbf{assert} e) &= \text{Rread}(e) \\
\text{REF}(\epsilon = e) &= \text{Lread}(\epsilon) \cup \text{Rread}(e)
\end{aligned}$$

$$\begin{aligned}
\text{DEF}(\tau \alpha = e) &= \text{prims}(\mathbf{Vvar} \alpha) \\
\text{DEF}(\tau \alpha = f(e_1, \dots)) &= \text{prims}(\mathbf{Vvar} \alpha) \cup \bigcup \text{Pointed}(e_i) \\
\text{DEF}(\tau \alpha = f(e_1, \dots) \mathbf{R} S_R \mathbf{W} S_W) &= \text{prims}(\mathbf{Vvar} \alpha) \cup S_W \\
\text{DEF}(\mathbf{assert} e) &= \emptyset \\
\text{DEF}(\epsilon = e) &= \text{prims}([\epsilon]_{L_0}^*)
\end{aligned}$$

Some simple facts about the \approx relation.

Fact 6 $\Omega \approx_{S_1} \Omega' \wedge \Omega \approx_{S_2} \Omega' \implies \Omega \approx_{S_1 \cup S_2} \Omega'$.

Fact 7 \approx_S is an equivalence relation.

Fact 8 $\Omega \approx_{S_1} \Omega' \wedge S_2 \subseteq S_1 \implies \Omega \approx_{S_2} \Omega'$.

This lemma says executing a command does not change values defined outside $\text{DEF}(c)$.

Lemma 17 $\forall \mu \notin \text{DEF}(c), (c, \Omega) \rightarrow_C (\text{cont}, \Omega', \beta) \implies \Omega \approx_{\{\mu\}} \Omega'$

D.1 Proof of Theorem 4

Here we restate the theorem:

If both $(c_1; \dots; c_{j-1}, \Omega) \rightarrow_C (\text{cont}, \Omega_{j-1}, \beta_{j-1})$ and $(c_j, \Omega_{j-1}) \rightarrow_C (\text{cont}, \Omega_j, \beta)$, then

$$\begin{aligned}
&\forall \Omega'_{j-1}, \Omega_{j-1} \approx_{\text{DEP}(-, j)} \Omega'_{j-1} \implies \\
&\exists \Omega'_j, (c_j, \Omega'_{j-1}) \rightarrow_C (\text{cont}, \Omega'_j, \beta) \wedge \Omega_j \approx_{\text{DEP}(j, -)} \Omega'_j.
\end{aligned}$$

This theorem indicates two facts about the live variable analysis:

1. The variables dependent by c_j , i.e., $\text{DEP}(-, j)$, are enough to execute c_j , and
2. When the variables in $\text{DEP}(-, j)$ are the same in the starting stacks, after executing c_j , the ending stacks are no different in terms of $\text{DEP}(j, -)$.

These two facts allow a program chunk to safely execute with part of the stack rather than reconstructing the entire stack.

Proof. Firstly, using lemma 3, we know Ω_{j-1} and Ω'_{j-1} are similar up to $\text{REF}(c_j)$. By the definition of $\text{REF}(c_j)$, we know during executing c_j , it only reads values from $\text{REF}(c_j)$. Thus, the behavior is also determined by them. We do not show detailed proof here.

So there must exist an ending stack Ω'_j where c_j can reach from Ω'_{j-1} and $\Omega_j \approx_{\text{DEF}(c_j)} \Omega'_j$. Now using lemma 17, we know $\Omega_{j-1} \approx_{\text{REF}(c_j)-\text{DEF}(c_j)} \Omega_j$ and $\Omega'_{j-1} \approx_{\text{REF}(c_j)-\text{DEF}(c_j)} \Omega'_j$. Since $\Omega_{j-1} \approx_{\text{REF}(c_j)} \Omega'_{j-1}$, we can use fact 7 and 8 to derive $\Omega_j \approx_{\text{REF}(c_j)-\text{DEF}(c_j)} \Omega'_j$.

Finally, fact 6 allows us to combine the similarity between Ω_j and Ω'_j in terms of $\text{DEF}(c_j)$ and $\text{REF}(c_j) - \text{DEF}(c_j)$ into $\Omega_j \approx_{\text{REF}(c_j) \cup \text{DEF}(c_j)} \Omega'_j$, then use lemma 2 we can prove $\Omega_j \approx_{\text{DEP}(j,-)} \Omega'_j$. ■

Theorem 18 *Using Theorem 4's setting, if $i < j < k$, then $\Omega_{j-1} \approx_{\text{DEP}(i,k)} \Omega_j$.*

Proof. We will prove it using lemma 17 by showing each variable in $\text{DEP}(i, k)$ does not occur in $\text{DEF}(c_j)$, i.e., $\text{DEF}(c_j) \cap \text{DEP}(i, k) = \emptyset$. By definition, if $\mu \in \text{DEP}(i, k)$, then $\mu \mapsto k \in \text{LIVE}(i)$. We then use the following lemma to prove $\mu \notin \text{DEF}(c_j)$.

Lemma 19 $\mu \mapsto k \in \text{LIVE}(i) \implies \forall j \in [i, k], \mu \mapsto k \in \text{LIVE}(j)$.

This lemma is direct from the definition of $\text{LIVE}(-)$: a dependency $\mu \mapsto k$ is added when analyzing c_k and will stay in the live set until μ is used in a command. Since it still exists in $\text{LIVE}(i)$, all live sets between i and k must not resolve it.

Returning to the theorem, we know $\mu \mapsto k \in \text{LIVE}(j - 1)$ using this lemma. So by definition of $\text{LIVE}(j - 1)$, $\mu \notin \text{DEF}(c_j)$. ■

D.2 Compromise in analysis

Note this section's discussion assumes the reader has read Sec. 5.2. One part of the efficient-cutting problem's goal is to minimize the sum of all chunk's external dependencies. However, in many cases, the cutting cost could be smaller. Consider

```

1 pvt int x = ...;
2 pvt int y = x;
3 pvt int z = x;

```

Suppose the cutting cost of one `pvt int` value is N . If `Lian` decides to partition L1 into one chunk, L2 and L3 into another chunk, then the OPB's encoding considers the cutting cost to be $2N$. However, since x is not updated between L2 and L3, L3 can use x 's local value without synchronization.

We could have used a more accurate encoding, but that will significantly increase the complexity and slow down the solver. So here we make a compromise by moving the dependency between L1 and L3 to L2 and L3. Now the dependency becomes local, so it no longer contributes to the total cutting cost.

We can rigorously prove in most cases, this approach does not reduce the total cutting cost in the encoding: If the program is cut into $(\langle 1 \rangle, \langle 2, 3 \rangle)$, $(\langle 1, 2 \rangle, \langle 3 \rangle)$ or $(\langle 1 \rangle, \langle 2 \rangle, \langle 3 \rangle)$, the total cost does not change (even if there are other commands involved). There is only one case that results in bigger cutting cost: $(\langle 1, 3 \rangle, \langle 2 \rangle)$. In other words, we will overestimate the total cost. This may not be a bad thing as now our encoding is prone to cut the program without breaking its original structure. Also, if the program has better performance by cutting 2 out, then usually the user can rewrite the program by moving commands that have close data dependencies together.

E Detailed Deep Simulation

Here are some unlisted semantics rules.

$$\text{dload}(\Omega, \epsilon) = \begin{cases} \llbracket \text{load}(\Omega, \mu), \mu \rrbracket_R & \text{when } \llbracket \epsilon \rrbracket_{L_1} = \mu \\ \epsilon & \text{otherwise} \end{cases}$$

$$\text{dstore}(\Omega, \epsilon, e) = \begin{cases} \text{store}(\Omega, \mu, \llbracket e \rrbracket_{R_1}) & \text{when } \llbracket \epsilon \rrbracket_{L_1} = \mu \\ \Omega & \text{otherwise} \end{cases}$$

$$\begin{aligned} \llbracket x \rrbracket_{L_1}^\Omega &= \text{find}(\Omega, x) \\ \llbracket \epsilon[e] \rrbracket_{L_1}^\Omega &= \llbracket \epsilon \rrbracket_{L_1}^\Omega \llbracket \llbracket e \rrbracket_{R_1}^\Omega \rrbracket \\ \llbracket \epsilon.l \rrbracket_{L_1}^\Omega &= \llbracket \epsilon \rrbracket_{L_1}^\Omega.l \\ \llbracket *e \rrbracket_{L_1}^\Omega &= \epsilon \quad \text{when } \llbracket e \rrbracket_{R_1}^\Omega = \& \epsilon \end{aligned}$$

$$\begin{aligned} \llbracket \epsilon \rrbracket_{R_1}^\Omega &= \text{dload}(\Omega, \llbracket \epsilon \rrbracket_{L_1}^\Omega) \\ \llbracket \& \epsilon \rrbracket_{R_1}^\Omega &= \& \llbracket \epsilon \rrbracket_{L_1}^\Omega \\ \llbracket \ell \{e_0, \dots, e_{n-1}\} \rrbracket_{R_1}^\Omega &= \ell \{ \llbracket e_0 \rrbracket_{R_1}^\Omega, \dots, \llbracket e_{n-1} \rrbracket_{R_1}^\Omega \} \\ \llbracket f(e_1, \dots, e_n) \rrbracket_{R_1}^\Omega &= \begin{cases} (f)_1(\llbracket e_1 \rrbracket_{R_1}^\Omega, \dots, \llbracket e_n \rrbracket_{R_1}^\Omega) & \text{when } f \in \text{dom}(_)_0 \\ f(\llbracket e_1 \rrbracket_{R_1}^\Omega, \dots, \llbracket e_n \rrbracket_{R_1}^\Omega) & \text{otherwise} \end{cases} \end{aligned}$$

$$\text{new}(x) = \alpha \quad \llbracket e \rrbracket_{R_1}^\Omega = e'$$

$$\frac{}{(\tau x = e, \Omega) \hookrightarrow_C (\text{cont}, \Omega[\alpha \mapsto \llbracket e' \rrbracket_{R_0}])}$$

$$\llbracket \epsilon \rrbracket_{L_1}^\Omega = \epsilon' \quad \llbracket e \rrbracket_{R_1}^\Omega = e'$$

$$\frac{}{(\epsilon = e, \Omega) \hookrightarrow_C (\text{cont}, \text{dstore}(\Omega, \epsilon', e'))}$$

$$\frac{}{(\mathbf{return} e, \Omega) \hookrightarrow_C (\text{ret } \llbracket e \rrbracket_{R_1}^\Omega, \Omega)}$$

$$\frac{}{(\mathbf{assert} e, \Omega) \hookrightarrow_C (\text{cont}, \Omega)}$$

$$\frac{}{(c_1, \Omega) \hookrightarrow_C (\text{ret } e, \Omega_1)}$$

$$(c_1; c_2, \Omega) \hookrightarrow_C (\text{ret } e, \Omega_1)$$

$$(c_1, \Omega) \hookrightarrow_C (\text{cont}, \Omega_1) \quad (c_2, \Omega_1) \hookrightarrow_C (r, \Omega_2)$$

$$(c_1; c_2, \Omega) \hookrightarrow_C (r, \Omega_2)$$

$$\llbracket e \rrbracket_{R_1}^\Omega = \ell \text{ true} \quad (c_1, \Omega) \hookrightarrow_C (r, \Omega')$$

$$\frac{}{(\mathbf{if} e \text{ then } c_1 \text{ else } c_2, \Omega) \hookrightarrow_C (r, \Omega')}$$

$$\llbracket e \rrbracket_{R_1}^\Omega = \ell \text{ false} \quad (c_2, \Omega) \hookrightarrow_C (r, \Omega')$$

$$\frac{}{(\mathbf{if} e \text{ then } c_1 \text{ else } c_2, \Omega) \rightsquigarrow_C (r, \Omega')}$$

$$\llbracket e \rrbracket_{R_1}^\Omega = \ell \text{ true}$$

$$(c; \mathbf{while} e \text{ do } c, \Omega) \hookrightarrow_C (r, \Omega')$$

$$\frac{}{(\mathbf{while} e \text{ do } c, \Omega) \hookrightarrow_C (r, \Omega')}$$

$$\begin{array}{c}
\frac{\llbracket e \rrbracket_{R_1}^\Omega = \ell \text{ false}}{\text{(while } e \text{ do } c, \Omega) \hookrightarrow_C (\text{cont}, \Omega)} \\
\text{new}(x) = \alpha \quad \eta \in \{\mathbf{normal}, \mathbf{atomic}, \mathbf{plocal1}, \mathbf{box1}\} \\
f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \\
\forall i \in [1, n], \llbracket e_i \rrbracket_{R_1}^\Omega = e'_i \wedge \text{new}(x_i) = \alpha_i \\
(c, \bigcup \{\alpha_i \mapsto \llbracket e'_i \rrbracket_{R_1} \} \cdot \Omega) \hookrightarrow_C (\text{ret } e, \omega \cdot \Omega') \\
\hline
(\tau x = f(e_1, \dots, e_n), \Omega) \hookrightarrow_C (\text{cont}, \Omega'[\alpha \mapsto \llbracket e \rrbracket_{R_1}]) \\
\text{new}(x) = \alpha \quad \eta \in \{\mathbf{plocal2}, \mathbf{box2}\} \\
f \mapsto \mathbf{internal} (\eta, (\tau_1 x_1, \dots, \tau_n x_n) \rightarrow \tau, c) \in \Lambda^u \\
\forall i \in [1, n], \llbracket e_i \rrbracket_{R_1}^\Omega = e'_i \\
\hline
(\tau x = f(e_1, \dots, e_n), \Omega) \hookrightarrow_C (\text{cont}, \Omega)
\end{array}$$

Lemma 20 $\llbracket e \rrbracket_R^\Omega = v \implies \exists e', \llbracket e \rrbracket_{R_1}^{\mathcal{Q}(\Omega)} = e' \wedge \llbracket e' \rrbracket_R^\Omega = v.$

Lemma 21 $\llbracket \epsilon \rrbracket_L^\Omega = \mu \implies \exists \epsilon', \llbracket \epsilon \rrbracket_{L_1}^{\mathcal{Q}(\Omega)} = \epsilon' \wedge \llbracket \epsilon' \rrbracket_L^\Omega = \mu.$

F Distribution Proofs

We denote $\text{chunk}^{-1}(t) := \{i \mid \text{chunk}(i) = t\}$. We let ϕ_j^t be the set of L-values that are used in $[1, j] \cap \text{chunk}^{-1}(t)$ and needed in $[j+1, n] \cap \text{chunk}^{-1}(t)$.

$$\phi_j^t := \bigcup_{\substack{i \leq j < k \\ \text{chunk}(i)=t \\ \text{chunk}(k)=t}} \text{DEP}(i, k)$$

Theorem 6 can be derived from the following lemma.

Lemma 22 $\forall j \in \text{chunk}^{-1}(t), \Omega_j^{\text{pre}} \approx_S \Omega_{j-1}$ where $S = \phi_{j-1}^t \cup \text{DEP}(-, j)$.

We prove it together with the following lemma by induction on the sequential program.

Lemma 23 $\forall j \in [1, n], \Omega_j^{\text{post}} \approx_S \Omega_j$ where $S = \phi_j^t$.

Subgoal 1. Suppose $\text{chunk}(j) = t$ and

$$\Omega_{j-1}^{\text{post}} \approx_S \Omega_{j-1} \text{ where } S = \phi_{j-1}^t,$$

we want to derive

$$\Omega_j^{\text{pre}} \approx_{S'} \Omega_{j-1} \text{ where } S' = \phi_{j-1}^t \cup \text{DEP}(j, -).$$

We know $\Omega_{j-1}^{\text{post}} \approx_S \Omega_j^{\text{pre}}$ as sync_j only updates external dependencies, so fact 7 ensures $\Omega_j^{\text{pre}} \approx_S \Omega_{j-1}$. Also, theorem 5 guarantees the correctness of sync_j , i.e.,

$$\Omega_j^{\text{pre}} \approx_{S''} \Omega_{j-1} \text{ where } S'' = \bigcup_{\text{chunk}(i) \neq t} \text{DEP}(i, j).$$

Finally, we can prove

$$\phi_{j-1}^t \cup \bigcup_{\text{chunk}(i) \neq t} \text{DEP}(i, j) \subseteq \phi_{j-1}^t \cup \text{DEP}(-, j),$$

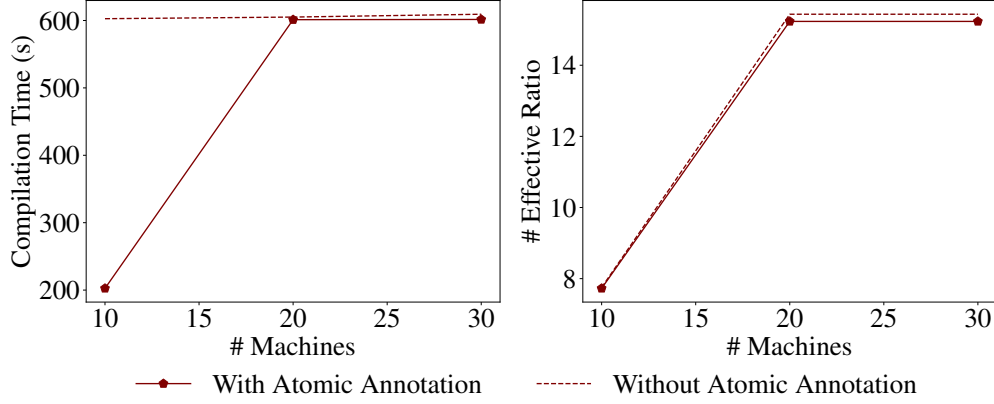


Figure 20: The effect of adding atomic annotations to the MT program for 16-block inputs.

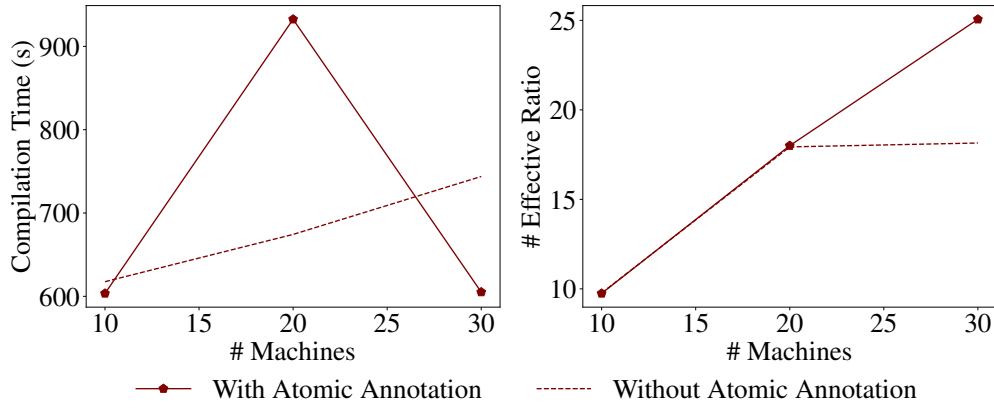


Figure 21: The effect of adding atomic annotations to the MT program for 64-block inputs.

then use fact 6 and 8 to derive $\Omega_j^{\text{pre}} \approx_{S'} \Omega_{j-1}$. ■

Subgoal 2. Suppose $\text{chunk}(j) = t$ and

$$\Omega_j^{\text{pre}} \approx_{S'} \Omega_{j-1} \text{ where } S' = \phi_{j-1}^t \cup \text{DEP}(-, j),$$

we want to derive

$$\Omega_j^{\text{post}} \approx_S \Omega_j \text{ where } S = \phi_j^t.$$

We divide ϕ_j^t into two parts S_1 and S_2 where $S_1 = \bigcup_{\text{chunk}(k)=t} \text{DEP}(j, k)$ contains all internal dependencies coming out of j , and $S_2 = \phi_j^t - S_1$ contains all internal dependencies coming before j . Theorem 4 ensures $\Omega_j^{\text{post}} \approx_{S_1} \Omega_j$, and theorem 18 ensures $\Omega_{j-1} \approx_{S_2} \Omega_j$ and $\Omega_j^{\text{pre}} \approx_{S_2} \Omega_j^{\text{post}}$. Notice $S_2 \subseteq \phi_{j-1}^t$, so we can combine with the hypothesis to derive $\Omega_j^{\text{post}} \approx_{S_2} \Omega_j$. Finally, we can use fact 6 to derive the goal. ■

G Supplementary Figures

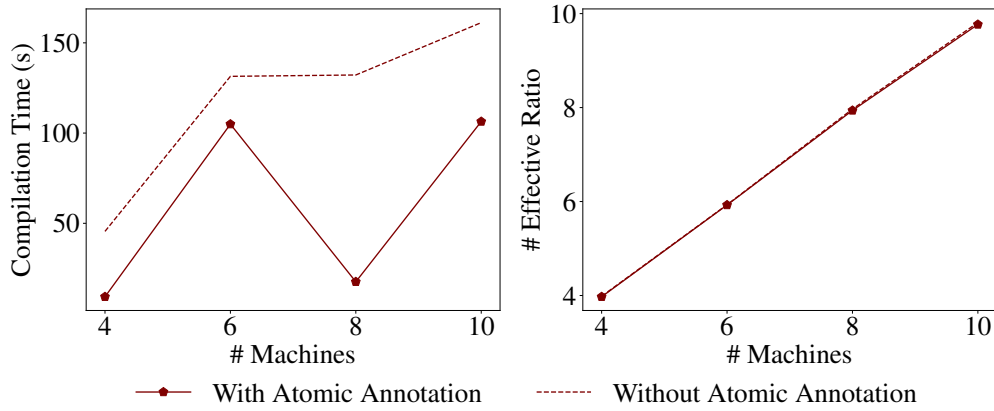
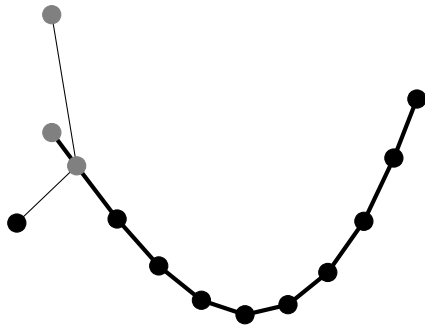
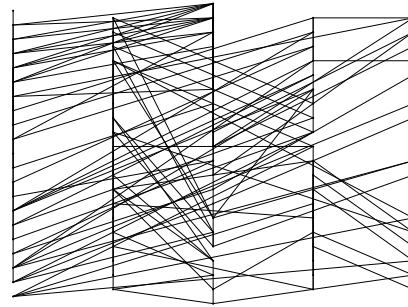


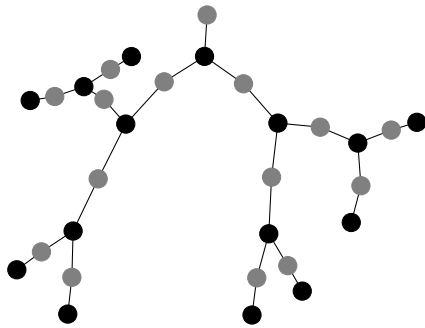
Figure 22: The effect of adding atomic annotations to the MT program for 8-block inputs.



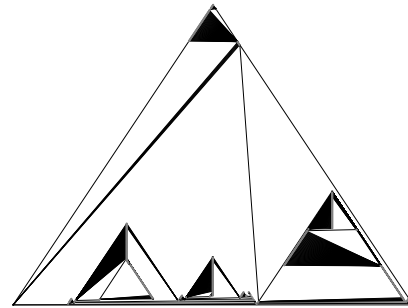
(a) Program structure of GD with `update_all()` annotated with atomic keyword



(b) Program structure of GD without `update_all()` being atomic



(a) Program structure of MT with `sha256_node()` annotated with atomic keyword



(b) Program structure of MT without `sha256_node()` being atomic